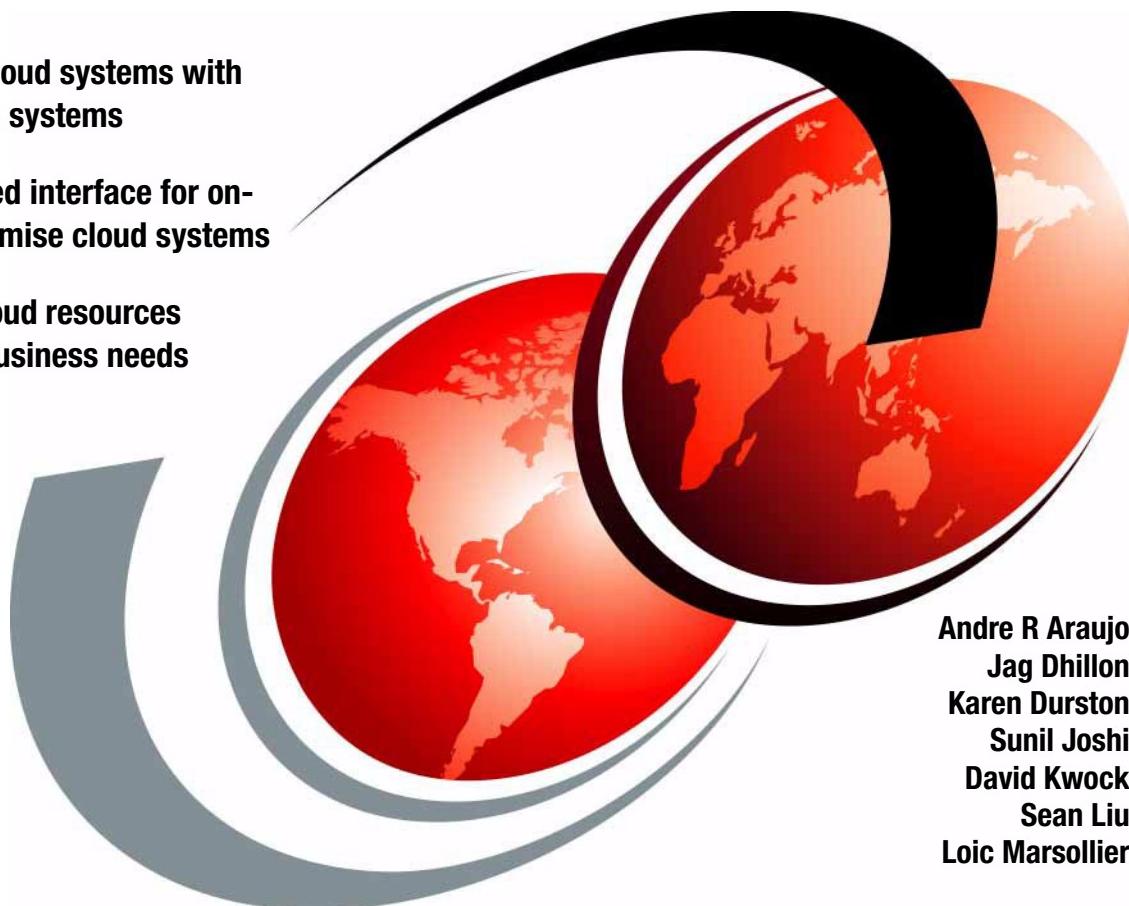


Hybrid Cloud Integration and Monitoring with IBM WebSphere Cast Iron

Integrate cloud systems with
on-premise systems

Use a unified interface for on-
and off-premise cloud systems

Allocate cloud resources
based on business needs



Andre R Araujo
Jag Dhillon
Karen Durston
Sunil Joshi
David Kwock
Sean Liu
Loic Marsollier

Redbooks



International Technical Support Organization

**Hybrid Cloud Integration and Monitoring with IBM
WebSphere Cast Iron**

December 2012

Note: Before using this information and the product it supports, read the information in "Notices" on page vii.

First Edition (December 2012)

This edition applies to IBM WebSphere Cast Iron V6.1, Service Management Extensions for Hybrid Cloud, IBM Tivoli Monitoring V6.2.2, IBM Tivoli Netcool V5.1.1, IBM Tivoli Service Automation Manager V7.2.2, IBM WebSphere ILOG JRules V7.1.1.

© Copyright International Business Machines Corporation 2012. All rights reserved.

Note to U.S. Government Users Restricted Rights -- Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Contents

| | |
|--|------|
| Notices | vii |
| Trademarks | viii |
| | |
| Preface | ix |
| The team who wrote this book | .ix |
| Now you can become a published author, too! | .xi |
| Comments welcome..... | .xi |
| Stay connected to IBM Redbooks | xii |
| | |
| Part 1. Introduction | 1 |
| | |
| Chapter 1. Overview of hybrid clouds | 3 |
| 1.1 Integration challenges to cloud adoption..... | 5 |
| 1.2 What is hybrid cloud | 6 |
| 1.3 IBM Hybrid Cloud solution..... | 8 |
| 1.4 Hybrid cloud dimensions | 12 |
| 1.5 Hybrid cloud scenarios | 12 |
| | |
| Chapter 2. Integrating cloud solutions with Cast Iron | 15 |
| 2.1 Overview of IBM WebSphere Cast Iron | 16 |
| 2.1.1 Appliance model | 17 |
| 2.1.2 Cast Iron Live | 19 |
| 2.2 Examples of Cast Iron implementations | 21 |
| 2.3 How WebSphere Cast Iron integrates applications..... | 22 |
| 2.4 Service Management Extensions for Hybrid Cloud..... | 25 |
| 2.4.1 Software components of the solution | 25 |
| 2.4.2 Prerequisites | 25 |
| 2.5 Usage scenarios | 26 |
| 2.5.1 New client scenario | 27 |
| 2.6 Clients with a WebSphere Cast Iron setup scenario | 27 |
| 2.6.1 Clients with a Tivoli Monitoring setup scenario | 28 |
| 2.7 Clients with a Tivoli Service Automation Manager setup scenario | 28 |
| 2.7.1 Clients with a LDAP v3-compliant Directory Server scenario | 29 |
| 2.8 Installing the Service Management Extensions for Hybrid Cloud | 29 |
| 2.8.1 Installing the extensions through the Cast Iron design console | 30 |
| 2.9 Example scenario | 32 |
| 2.9.1 Provisioning..... | 32 |
| 2.9.2 Monitoring | 35 |
| 2.9.3 Directory Synchronization | 37 |

| | |
|---|-----------|
| 2.9.4 Workload Governance and Management | 38 |
| 2.10 Summary | 40 |
| Chapter 3. IBM Cloud Service Management Platform | 41 |
| 3.1 Cloud computing architecture | 42 |
| 3.1.1 Key roles in a cloud ecosystem. | 42 |
| 3.1.2 Common Cloud Management Platform. | 43 |
| 3.2 Mapping IBM products to the CCMP. | 47 |
| 3.2.1 Solutions supporting the BSS | 48 |
| 3.2.2 Solutions supporting the OSS. | 49 |
| 3.3 Cloud Service Management Platform for Hybrid Cloud. | 55 |
| 3.3.1 Updated components to support Hybrid Cloud | 56 |
| 3.3.2 New components to support Hybrid Cloud | 57 |
| Part 2. Usage scenarios | 59 |
| Chapter 4. Monitoring scenario | 61 |
| 4.1 Monitoring scenario overview | 62 |
| 4.2 Installation and configuration. | 64 |
| 4.2.1 Install Service Management Extensions for Hybrid Cloud plug-ins. | 64 |
| 4.2.2 Install IBM Cloud Gateway Agent on IBM SmartCloud Enterprise | 68 |
| 4.2.3 Creating the integration. | 75 |
| 4.3 Verifying the installation | 81 |
| 4.3.1 Log in to IBM Tivoli Monitoring | 81 |
| 4.3.2 Verifying the new Linux agent is shown in dashboard | 83 |
| 4.3.3 Verifying the monitoring scenario is working. | 84 |
| 4.4 Running a CPU monitoring test for the Workload Governance and Management scenario | 88 |
| 4.5 Lessons learned | 89 |
| 4.6 Preferred practices | 91 |
| 4.7 Summary | 92 |
| Chapter 5. Provisioning scenario | 93 |
| 5.1 Provisioning scenario overview | 95 |
| 5.2 Installing and configuring the Provisioning scenario | 96 |
| 5.2.1 Installing the Service Management Extensions for Hybrid Cloud | 96 |
| 5.2.2 Creating the integration. | 96 |
| 5.2.3 Installing the Tivoli Service Automation Manager Hybrid Cloud Extension | 102 |
| 5.2.4 Configure Tivoli Service Automation Manager for a hybrid cloud | 123 |
| 5.3 Scenario usage | 138 |
| 5.3.1 Creating a project with servers in the IBM SmartCloud Enterprise | 139 |
| 5.3.2 Adding one or multiple servers to a project. | 145 |
| 5.3.3 Removing a server from a project. | 147 |

| | |
|---|------------|
| 5.3.4 Cancelling a project | 149 |
| 5.3.5 Logging in to the Tivoli Service Automation Manager | 150 |
| 5.3.6 Logging in to the Maximo Start Center | 150 |
| 5.4 Lessons learned | 150 |
| 5.5 Summary | 150 |
| Chapter 6. Directory Synchronization scenario | 153 |
| 6.1 Scenario overview | 154 |
| 6.2 Implementing the directory synchronization | 156 |
| 6.2.1 Install Service Management Extensions for Hybrid Cloud plug-ins | 157 |
| 6.2.2 Create directory entries for WebSphere Cast Iron integration | 159 |
| 6.2.3 Create the WebSphere Cast Iron Directory sync integration | 162 |
| 6.3 Running and testing the scenario | 168 |
| 6.4 How the integration works | 170 |
| 6.5 Lessons learned | 171 |
| 6.6 Summary | 172 |
| Chapter 7. Workload Governance and Management scenario | 173 |
| 7.1 Scenario overview | 174 |
| 7.1.1 Tivoli Service Automation Manager | 175 |
| 7.1.2 ILOG JRules Execution Server | 175 |
| 7.1.3 Tivoli Monitoring | 176 |
| 7.1.4 Netcool Impact | 176 |
| 7.2 Installation and configuration | 176 |
| 7.2.1 Prerequisites and installation | 176 |
| 7.2.2 TSAM additional configuration | 178 |
| 7.2.3 Deploying the Event Correlation Service on the Tivoli Netcool/Impact server | 179 |
| 7.2.4 Editing the Event Correlation Service configuration file | 179 |
| 7.2.5 Installing the WebSphere ILOG JRules Rules Execution Server | 180 |
| 7.2.6 Verifying the installation of the Rules Execution Server | 181 |
| 7.2.7 Deploying the hybrid cloud RuleApp | 181 |
| 7.2.8 Additional Tivoli Monitoring configuration | 181 |
| 7.3 Scenario walk-through | 182 |
| 7.3.1 Resources added to cloud project to support business policy | 183 |
| 7.3.2 Resources decommissioned from the cloud project to support business policy | 192 |
| 7.4 Lessons learned | 199 |
| 7.5 Summary | 200 |
| Related publications | 201 |
| IBM Redbooks | 201 |
| Online resources | 201 |
| Help from IBM | 203 |

Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing, IBM Corporation, North Castle Drive, Armonk, NY 10504-1785 U.S.A.

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this IBM product and use of those websites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming

techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs.

Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. These and other IBM trademarked terms are marked on their first occurrence in this information with the appropriate symbol (® or ™), indicating US registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the Web at <http://www.ibm.com/legal/copytrade.shtml>

The following terms are trademarks of the International Business Machines Corporation in the United States, other countries, or both:

| | | |
|-------------------------|---------------------|---|
| Cast Iron® | LotusLive™ | Redpaper™ |
| DataPower® | Lotus® | Redbooks (logo)  ® |
| Dynamic Infrastructure® | Maximo® | Service Request Manager® |
| i5/OS™ | Netcool® | Smarter Cities® |
| IBM SmartCloud™ | Passport Advantage® | Tivoli® |
| IBM® | Power Systems™ | WebSphere® |
| ILOG® | Redbooks® | |

The following terms are trademarks of other companies:

Linux is a trademark of Linus Torvalds in the United States, other countries, or both.

Microsoft, Windows, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

Java, and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Other company, product, or service names may be trademarks or service marks of others.

Preface

Cloud computing is fast emerging as a new consumption and delivery model for IT solutions. Two distinct cloud patterns have evolved. Private clouds are custom cloud applications that are on premise to companies. Public clouds are owned and managed by a service provider and are delivered on a pay-per-use basis. Although the debate between the use of private versus public cloud continues, the industry is moving ever faster toward a hybrid solution. Hybrid clouds address the integration of traditional IT environments with one or more clouds, private and public.

IBM® WebSphere® Cast Iron® along with the IBM Tivoli® Service Management Extensions for Hybrid Cloud supports the integration of private and public clouds. The Service Management Extensions for Hybrid Cloud supports the following scenarios:

- ▶ Provisioning hybrid cloud resources
- ▶ Monitoring hybrid cloud resources
- ▶ Governing and managing workloads that use hybrid cloud resources
- ▶ Provisioning users from LDAP to IBM LotusLive™

This IBM Redbooks® publication is intended for application integrators, integration designers, and administrators evaluating or using IBM WebSphere Cast Iron. In addition, executives, business leaders, and architects who are looking for a way to integrate cloud applications with their on-premise applications are shown how WebSphere Cast Iron can help to resolve their integration challenges. The book helps you gain an understanding of hybrid cloud use cases, and explains how to integrate cloud and on-premise applications quickly and simply.

The team who wrote this book

This book was produced by a team of specialists from around the world working at the International Technical Support Organization, San Jose Center.

Andre R Araujo is a System Management Information Technology Architect working with Tivoli Management Product in Brazil. He has fifteen years of experience in servers and systems support. Andre is an IBM Tivoli Certified Advanced Deployment Professional. He holds a degree in Telecommunication Engineering and has an MBA in Network Computer Management. His areas of expertise include UNIX/Linux support, Cisco networking, networking security,

infrastructure and networking management. He has contributed to IBM Redbooks publications about Tivoli and IBM Netcool® products.

Jag Dhillon is a WebSphere Channels Technical Professional, IBM Australia. He has more than 15 years of experience in the Systems Infrastructure field. For the last seven years he has worked as an IBM Consultant on projects involving middleware applications. Jag holds technical certifications from various vendors, and has delivered a number of official IBM Education classes in Australia and New Zealand over the past few years.

Karen Durston is a course developer and instructor with WebSphere Education, and previously worked with Tivoli Education. She is an Advisory Software Engineer with 15 years of experience in technical enablement between the two brands. Karen holds a Masters of Computer Science degree from the University of Illinois at Urbana-Champaign and a B.A. with High Honors from the College of William and Mary in Virginia. Her areas of expertise include WebSphere Cast Iron, ITCAM for Application Diagnostics, and ITCAM for SOA.

Sunil Joshi is a Senior Enterprise Architect in the U.S. He has been with IBM for 11 years. His areas of expertise include Enterprise Architecture, Cloud Computing, Business Analytics, and IT delivery methods. Sunil has written extensively about cloud computing, methodologies, and crowdsourcing, and has presented at several forums such as the IBM Technical Leadership Exchanges and IBM Academy of Technology conferences.

David Kwock is a Service Management Solution Architect in the U.S. He has twelve years of experience in the IT Service Management field, and has worked at IBM for five years. His areas of expertise include Service Management, Service-Oriented Architecture, Cloud Computing and IBM Dynamic Infrastructure®. David has written extensively about IT Service Management, Service Oriented Architecture and Cloud Computing, and has contributed to an IBM Redpaper™ publication about Tivoli Provisioning Manager for OS deployment in a retail environment and to an IBM Redbooks publication about integrating Tivoli solutions. He has also presented on these topics at several conferences, including IBM Executive Summits and Gartner conferences.

Sean Liu is a Senior Systems and Storage Consultant with IBM China. He has 16 years of experience in the areas of Information Technology and Enterprise Datacenter environments. Sean works for IBM Growth Markets Unit Advanced Technical Skills, supporting IBM Power Systems™ Technical sales team activities related to Power Systems, Storage Subsystems, Virtualization, Business Continuity, Disaster Recovery, Cloud, and Next Generation Platform.

Loic Marsollier is a Cloud Computing IT Architect with 32 years of experience. In 1998, he became certified in Production and Inventory Management (CPIM). He joined IBM in 2001, and most recently has been working with the IBM PSSC

Customer Center in Montpellier, France. This center of excellence provides comprehensive customer briefings, architecture design, benchmarks, testing and training facilities to demonstrate benefits of IBM systems, software and solutions such as Cloud Computing and IBM Smarter Cities®. Loic spent five years as an IT Specialist with WebSphere, and then worked as a Software IT Architect, supporting the entire IBM software portfolio in the retail industry. He has received IBM IT Specialist certification and IT Architect certification.

Thanks to the following people for their contributions to this project:

Chris Rayns
International Technical Support Organization, Poughkeepsie Center

Carla Sadtler
International Technical Support Organization, Raleigh Center

Now you can become a published author, too!

Here's an opportunity to spotlight your skills, grow your career, and become a published author—all at the same time! Join an ITSO residency project and help write a book in your area of expertise, while honing your experience using leading-edge technologies. Your efforts will help to increase product acceptance and customer satisfaction, as you expand your network of technical contacts and relationships. Residencies run from two to six weeks in length, and you can participate either in person or as a remote resident working from your home base.

Find out more about the residency program, browse the residency index, and apply online at:

ibm.com/redbooks/residencies.html

Comments welcome

Your comments are important to us!

We want our books to be as helpful as possible. Send us your comments about this book or other IBM Redbooks publications in one of the following ways:

- ▶ Use the online **Contact us** review Redbooks form found at:

ibm.com/redbooks

- ▶ Send your comments in an email to:
redbooks@us.ibm.com
- ▶ Mail your comments to:
IBM Corporation, International Technical Support Organization
Dept. HYTD Mail Station P099
2455 South Road
Poughkeepsie, NY 12601-5400

Stay connected to IBM Redbooks

- ▶ Find us on Facebook:
<http://www.facebook.com/IBMRibooks>
- ▶ Follow us on Twitter:
<http://twitter.com/ibmredbooks>
- ▶ Look for us on LinkedIn:
<http://www.linkedin.com/groups?home=&gid=2130806>
- ▶ Explore new Redbooks publications, residencies, and workshops with the IBM Redbooks weekly newsletter:
<https://www.redbooks.ibm.com/Redbooks.nsf/subscribe?OpenForm>
- ▶ Stay current on recent Redbooks publications with RSS Feeds:
<http://www.redbooks.ibm.com/rss.html>



Part 1

Introduction

In this part we introduce the concept of hybrid clouds and the challenges to cloud adoption. We provide an overview of IBM WebSphere Cast Iron, and discuss how this product can play a key role in the integration of applications across public and private clouds.

This part also introduces a scenario that is used in Part 2, “Usage scenarios” on page 59 to illustrate the management capabilities available when using WebSphere Cast Iron as an integration solution.

The following chapters are included:

- ▶ Chapter 1, “Overview of hybrid clouds” on page 3
- ▶ Chapter 2, “Integrating cloud solutions with Cast Iron” on page 15
- ▶ Chapter 3, “IBM Cloud Service Management Platform” on page 41



Overview of hybrid clouds

In this chapter we provide an overview of hybrid clouds and their usage. The benefits for enterprises to implement cloud capabilities have been well documented: cost savings, enhanced productivity, and rapid deployment of new functionalities, to name a few.

Enterprises can use private clouds to scale requirements from lines of businesses quickly and in a cost-effective manner. The alternative to private clouds is the use of public clouds that are managed and delivered by service providers. Both cloud solutions have their advantages and challenges.

Adopting a cloud solution is not simply a smart thing for a business, but is becoming a necessity because that is where the marketplace is moving.

Cloud services are essentially delivered in three layers:

- ▶ Software as a Service (SaaS Applications): Use of applications provided and supported by vendors over the network
- ▶ Platform as a Service (PaaS): Use of vendor-managed and supported platforms to deploy custom or packaged cloud applications
- ▶ Infrastructure as a Service (IaaS): Use of basic infrastructure capabilities such as storage, network, CPU, and memory that is managed and supported by a vendor

The cloud concept allows businesses to quickly scale IT capabilities to meet a potential opportunity. The ease and speed of delivery is especially appealing to

large enterprises, where a centralized IT department usually takes disproportionate amounts of time to put together a solution.

Thus, many enterprises find themselves in the situation of having some form of both private and public clouds within their overall solution. In this chapter, we provide an overview of hybrid clouds and their usage.

In this chapter we discuss the following topics:

- ▶ Integration challenges to cloud adoption
- ▶ What is hybrid cloud
- ▶ IBM hybrid cloud solutions
- ▶ Hybrid cloud dimensions
- ▶ Hybrid cloud scenarios

1.1 Integration challenges to cloud adoption

One of the basic challenges for cloud adoption has been the struggle to integrate between the layers of cloud services and the enterprise. The following issues are some of the major aspects of the integration in a cloud landscape.

- ▶ Control issues:
 - Lack of centralized control over resources owned and allocated, including applying centralized security standards across private and public domains
 - Lack of a seamless solution to policy based allocation and decommissioning of resources
 - Creation of custom (SaaS) applications in silos that results in:
 - No “big picture” of the enterprise
 - Component reusability opportunities lost
 - Lack of a standardized architecture across the enterprise
- ▶ Standards and compliance issues that include:
 - Safeguarding confidential data
 - Federal standards and privacy law compliance
- ▶ Application issues that include:
 - Lack of a comprehensive data synchronization solution
 - Application management issues with version control, upgrades, and deploying new functionality

In summary, to fully exploit advantages of the cloud, enterprises must be able to:

- ▶ Integrate solutions across private and public cloud domains to scale up or scale down.
- ▶ Integrate on-premise and off-premise applications.
- ▶ Monitor and manage resources in a standardized way across on-premise and off-premise resources.
- ▶ Apply resource management governance policies in a consistent manner across on-premise and off-premise resources.

Those needs can be met with a *hybrid* approach to cloud adoption.

1.2 What is hybrid cloud

A hybrid solution bridges the gap between the on-premise existing systems and cloud applications, platform, and infrastructure, providing a rapid and easy-to-use method of setting up integrated solutions (Figure 1-1).

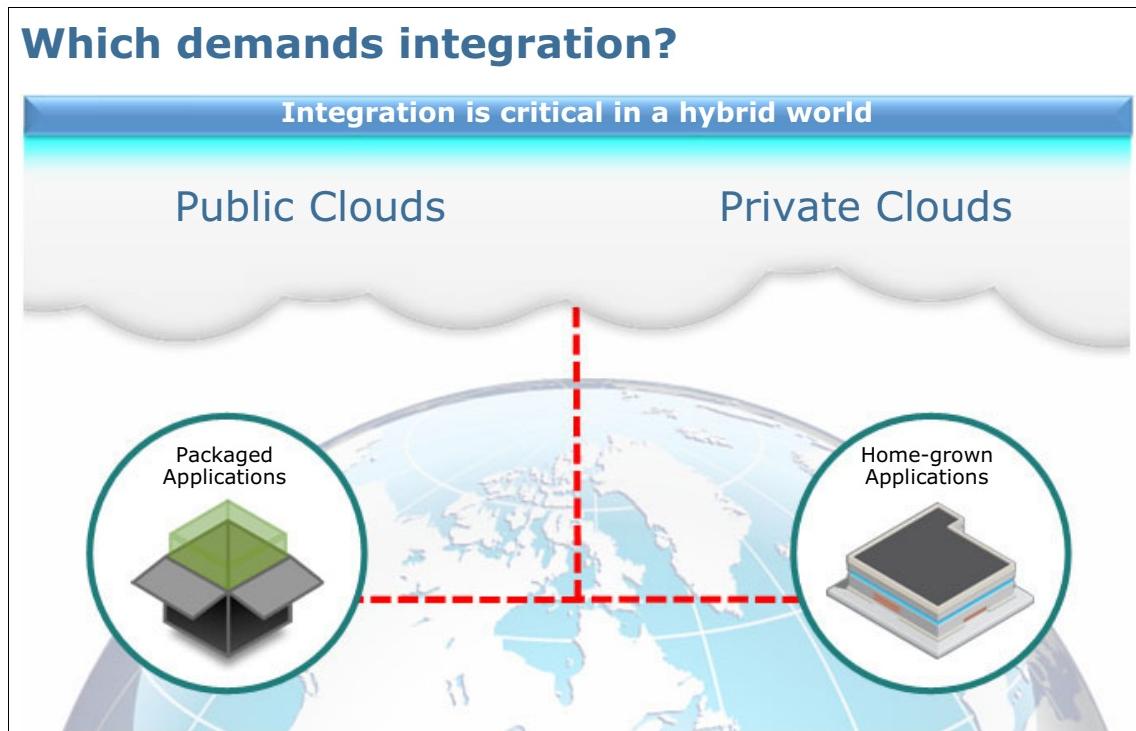


Figure 1-1 Hybrid cloud

The hybrid cloud theme addresses the integration of traditional IT environments with one or more clouds, on-premise and off-premise. The individual clouds and the traditional IT environments remain unique entities but are bound together by processes and technology. A fully comprehensive hybrid solution provides

integration capabilities at all three cloud services layers, as illustrated in Figure 1-2.

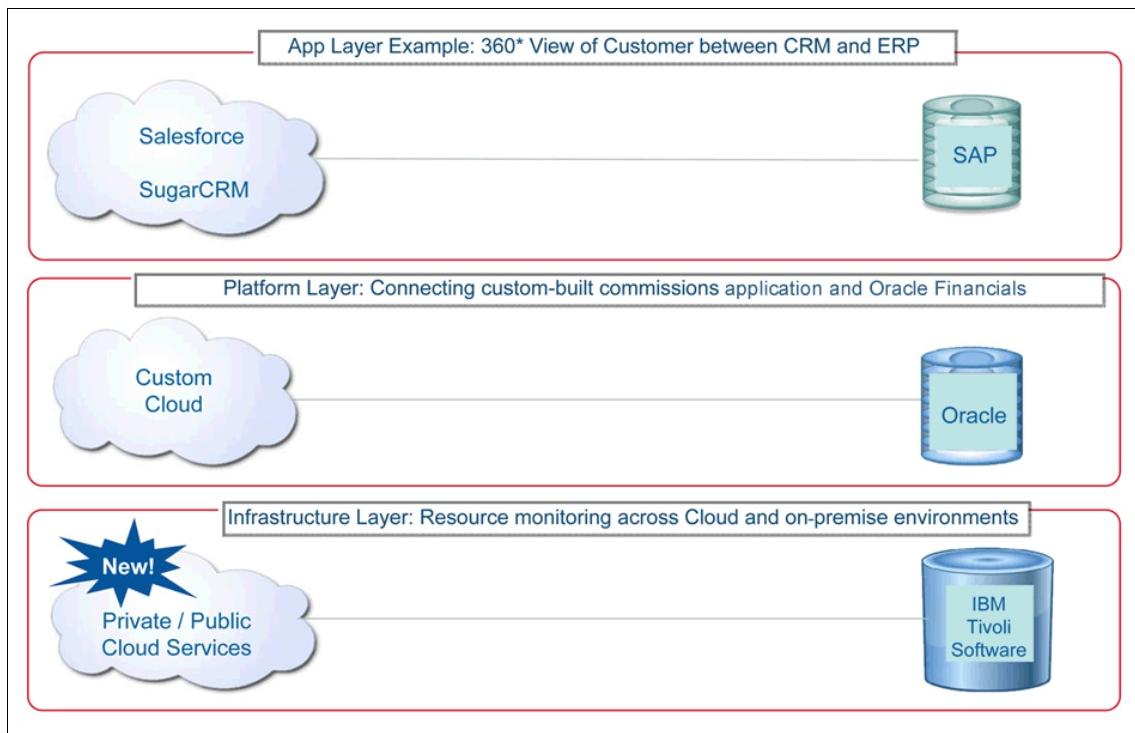


Figure 1-2 Three cloud services layers

By definition, a hybrid cloud solution provides:

- ▶ Integration of traditional on-premise IT environments with one or more clouds, either on-premise (private) clouds or off-premise (public) clouds
- ▶ Composition of two or more clouds that remain unique entities but are bound together by processes and technology to create a unique integrated environment, providing users a seamless experience across the individual clouds

1.3 IBM Hybrid Cloud solution

Before we delve into the IBM Hybrid Cloud solution, we examine the IBM Cloud Computing Reference Architecture (CCRA). The CCRA (Figure 1-3) covers the entire spectrum of cloud services, including integration capabilities.

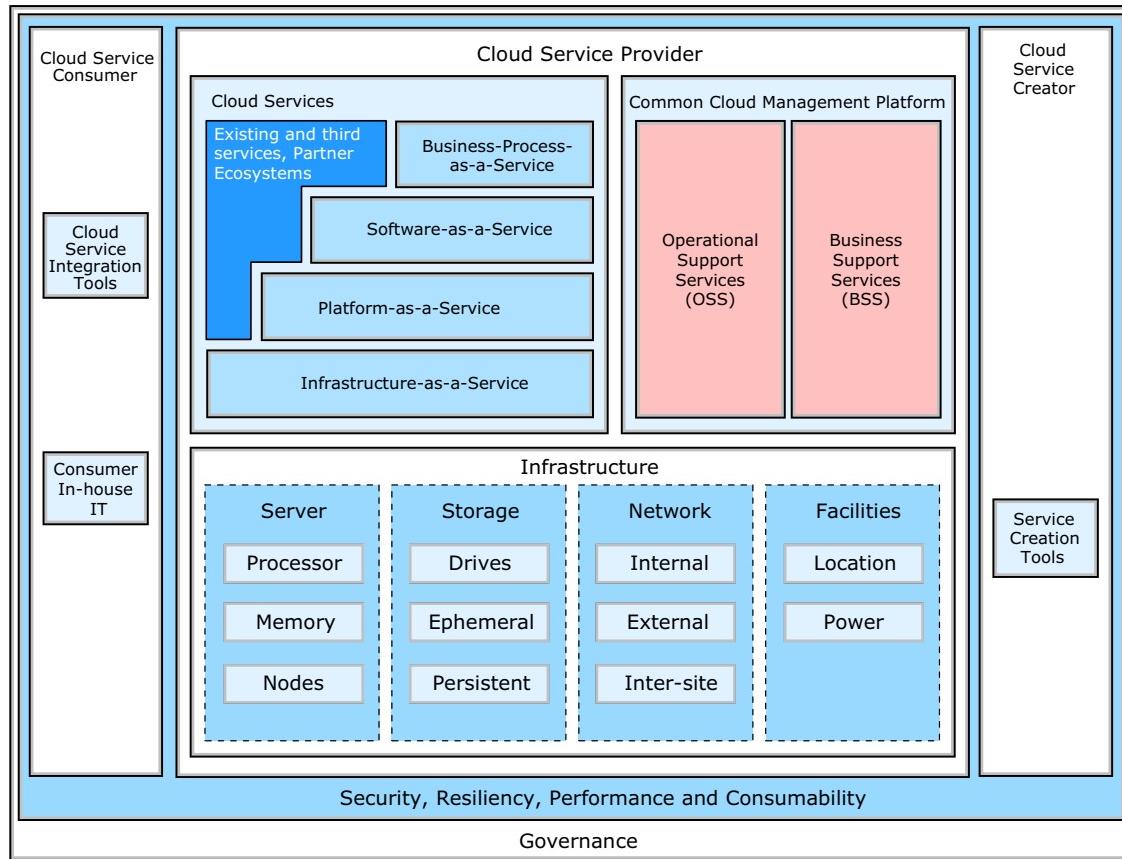


Figure 1-3 IBM Cloud Computing Reference Architecture (CCRA)

The IBM Hybrid Cloud solution consolidates and simplifies the integration and management of on-premise and cloud computing resources. The IBM Hybrid Cloud solution includes:

- ▶ WebSphere Cast Iron Software
 - This software enables companies to seamlessly integrate on-premise and public clouds. Enhancements to this platform include a complete multi-tenant cloud service that clients can use to design, run, and manage

all cloud integrations, and a new virtual appliance that can be installed on existing servers.

- ▶ Service Management Extensions for Hybrid Cloud
 - This helps clients to monitor and secure the management of resources on both public and private clouds. With Tivoli software for the Hybrid Cloud solution, clients can now manage off-premise resources to the same standard, and with the same infrastructure, as resources inside enterprise walls.

This Hybrid Cloud solution (Figure 1-4) provides visibility into off-premise clouds using the same user interface (UI) as those on premise, and access to applications that are enabled across vendors. The result is increased client confidence in the management of off-premise resources.

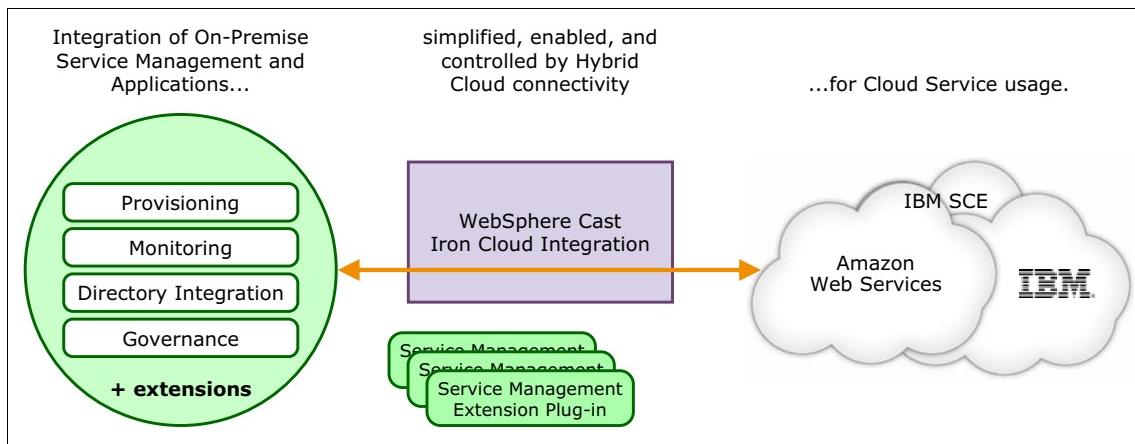


Figure 1-4 IBM Hybrid Cloud solution

The IBM Hybrid Cloud solution provides application integration capabilities, allowing clients to integrate application data from an off-premise service site into an on-premise application (Figure 1-5). As a result, clients can access public application data using the rapid integration capability of the WebSphere Cast Iron appliance.

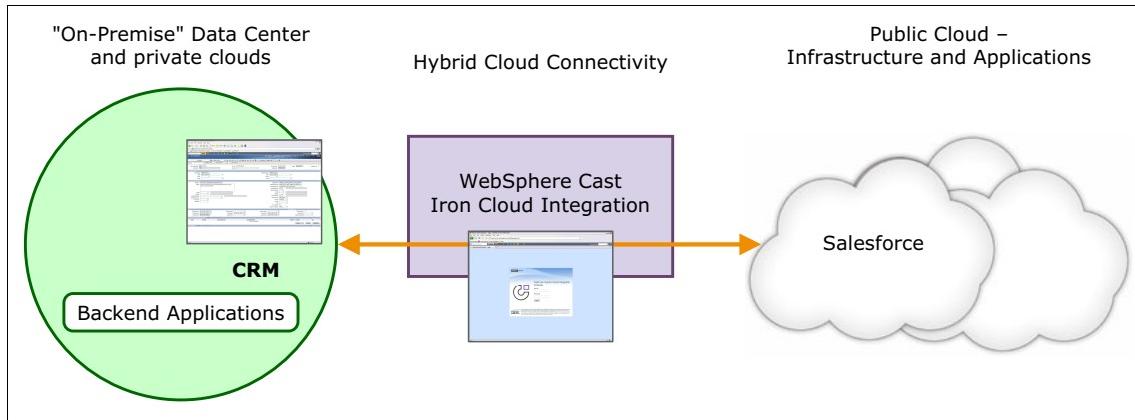


Figure 1-5 Application integration with the IBM Hybrid Cloud solution

The IBM Hybrid Cloud solution provides monitoring capabilities that allow clients to begin monitoring off-premise resources using their existing on-premise monitoring installation (Figure 1-6). As a result clients gain visibility into their off-premise “cloud” infrastructure using the same ITM monitoring solution as those on premise.

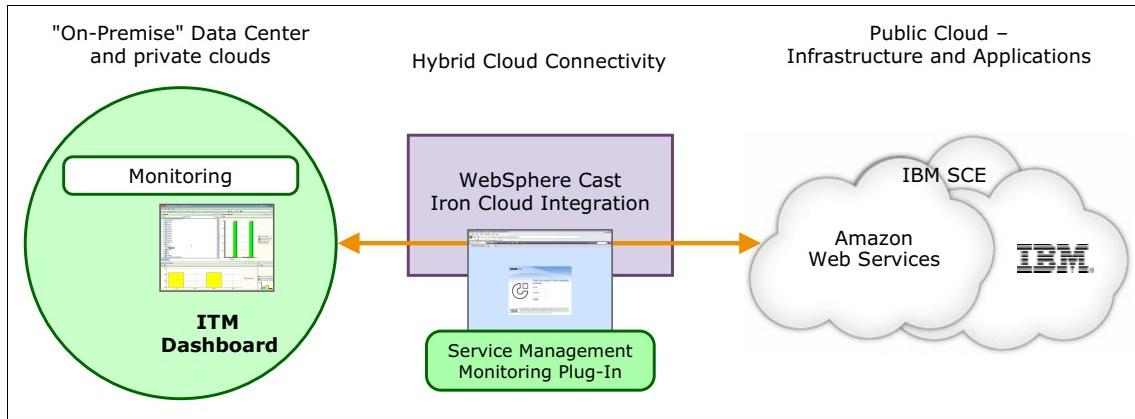


Figure 1-6 Solution monitoring with the IBM Hybrid Cloud solution

With the provisioning capabilities of the IBM Hybrid Cloud solution (Figure 1-7 on page 11), clients can provision off-premise resources using their existing

on-premise service automation installation. As a result, they extend service request automation into their off-premise cloud infrastructure using the same TSAM service automation management solution as those on premise.

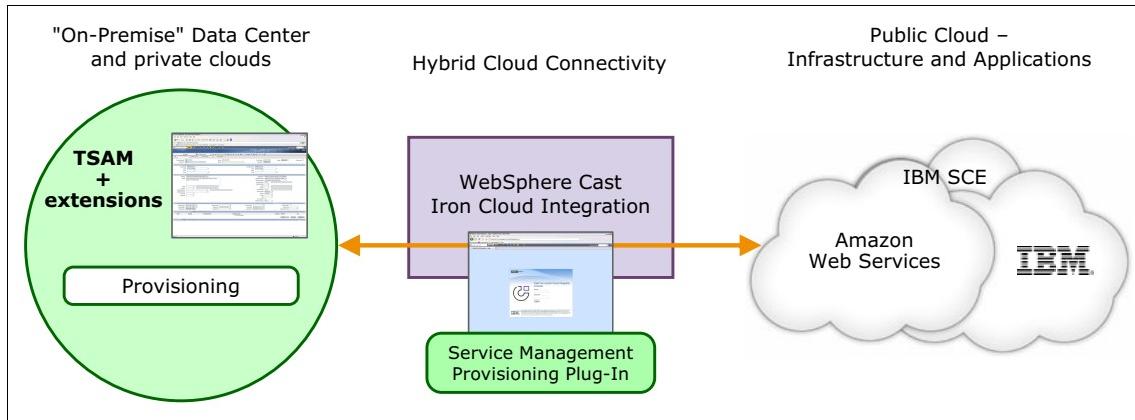


Figure 1-7 Solution provisioning with the IBM Hybrid Cloud solution

The governance capabilities of the Hybrid Cloud solution (Figure 1-8) allow an enterprise to use off-premise resources with its existing on-premise workload optimization installation. The result is an extension of the off-premise cloud infrastructure for workload optimization using the same Tivoli ITM, Omnibus, and TSAM solution as those on premise.

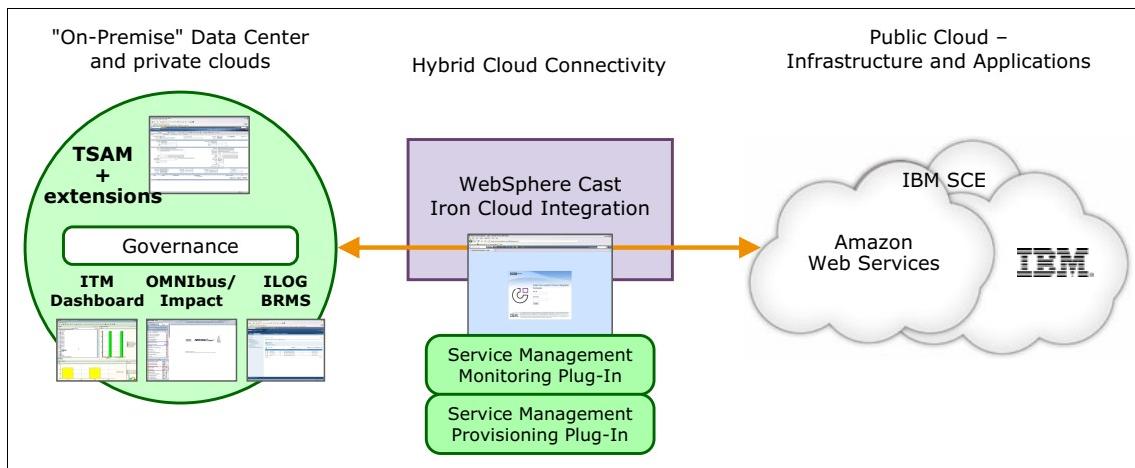


Figure 1-8 Solution governance with the IBM Hybrid Cloud solution

In the subsequent chapters, we describe in detail how the IBM WebSphere Cast Iron product along with the IBM Tivoli Service Management Extensions can provide a simplified yet comprehensive Hybrid Cloud solution.

1.4 Hybrid cloud dimensions

The hybrid cloud concept can be discussed within five broad dimensions.

- ▶ Application Integration (services and data): This dimension is concerned with connecting and integrating on-premise applications with off-premise services and data for spill-over capacity, functionality, and business data mapping.
- ▶ Security Integration: This dimension is concerned with integrating on-premise and off-premise identities, policies, and auditing systems. It includes ensuring proper security of off-premise cloud workload and securing management and payload interactions.
- ▶ Monitoring and aggregation: This dimension is concerned with the integration of sensing and monitoring off-premise cloud-delivered services (including infrastructure, platform, and applications) with on-premise aggregation and management systems. This dimension is also concerned with enablement of on-premise sensing services such as monitoring, metering, and event infrastructure to extend their reach into off-premise cloud-based services.
- ▶ Control and management: This dimension is concerned with controlling and managing access, availability, and enablement of service capacity in the off-premise cloud.
- ▶ Workload governance: This dimension is concerned with governing service request management of on-premise and off-premise services and resources, including managing access permission to cloud services.

1.5 Hybrid cloud scenarios

Typically, a hybrid cloud is the combination of public and private clouds that interoperate, allowing clients to outsource selected information and processing to the public cloud while keeping private data in their control. On-premise private clouds store and secure private and confidential data behind the enterprise's firewall. Public cloud capabilities are enabled through the Internet, without the need to acquire and maintain additional hardware.

There are also scenarios where public or private clouds are integrated with resources deployed and managed in traditional environments. A typical example is a hybrid web hosting scenario where the web server is in a public cloud and

the database server is in a private cloud or in a traditionally managed on-premise IT environment.

Another hybrid cloud scenario consists of cloud applications running in a public cloud and cooperating with partner applications in private clouds of different enterprises, enabling, for example, cloud-based supply chain solutions.

The following specific business scenarios fit a hybrid cloud solution:

- ▶ A retail company plans to move an order processing application onto its new private cloud, but the sales team has used Salesforce.com for some time. CIOs want to understand their role in delivering, managing, securing, and providing business resilience for these new services and service delivery technologies.
- ▶ A large finance company is thinking about building heterogeneous clouds, for example running Dev/Test x86 workloads on an open source hypervisor such as KVM, production x86 workloads on VMWare, and with ERP applications running on a Power-based system. The company wants the ability to manage and monitor these heterogeneous workloads in a single glass pane view.
- ▶ A large telecommunication company anticipates extreme spikes in workload on several of their business services during a new model launch. They anticipate several hundred thousand clients to consume business services for advocacy, but the volumes taper down soon after the launch.
- ▶ A retail company based in the southeast region of United States acquires a retail chain servicing the west coast. The company is now challenged to spin off integration projects for consolidation, and requires additional short-term hardware and software resources. These resources are thrown away after the integration projects are complete.
- ▶ The IT department of a retail company wants to dynamically define thresholds for determining workload overflow situations. The department wants a policy-based spillover capacity, by setting policies determining situations that qualify the use of an off-premise public cloud under qualified spillover situations.
- ▶ The IT department mentioned in the preceding scenario also has workload processing that requires access to on-premise data. The data in the on-premise repositories includes both business-critical, sensitive data and non-critical, non-sensitive data. Business-critical data must remain on premise and cannot be replicated to or processed on off-premise resources. When workload is to be processed off-premise, only non-critical data is to be replicated and made available on the off-premise public cloud.



Integrating cloud solutions with Cast Iron

IBM WebSphere Cast Iron Cloud Integration enables companies to integrate applications, regardless of whether the applications are located on premise or in public or private clouds. Now it also includes Service Management Extensions for Hybrid Cloud, which allows companies to integrate the monitoring, provisioning, events handling, and automating of service environments located both on and off premises.

In this chapter we provide an overview of WebSphere Cast Iron, and the architectural overview of an integrated Hybrid Cloud solution and the usage scenarios through the Service Management Extensions for Hybrid Cloud.

In this chapter we discuss the following topics:

- ▶ An overview of IBM WebSphere Cast Iron
- ▶ Examples of Cast Iron implementations
- ▶ Architectural overview of the integrated solution
- ▶ Installation of the Service Management Extensions for Hybrid Cloud
- ▶ Usage scenarios

2.1 Overview of IBM WebSphere Cast Iron

IBM WebSphere Cast Iron provides a solution that meets the challenge of integrating cloud applications with on-premise systems, cloud applications-to-cloud applications, and on-premise-to-on-premise applications, as illustrated in Figure 2-1.

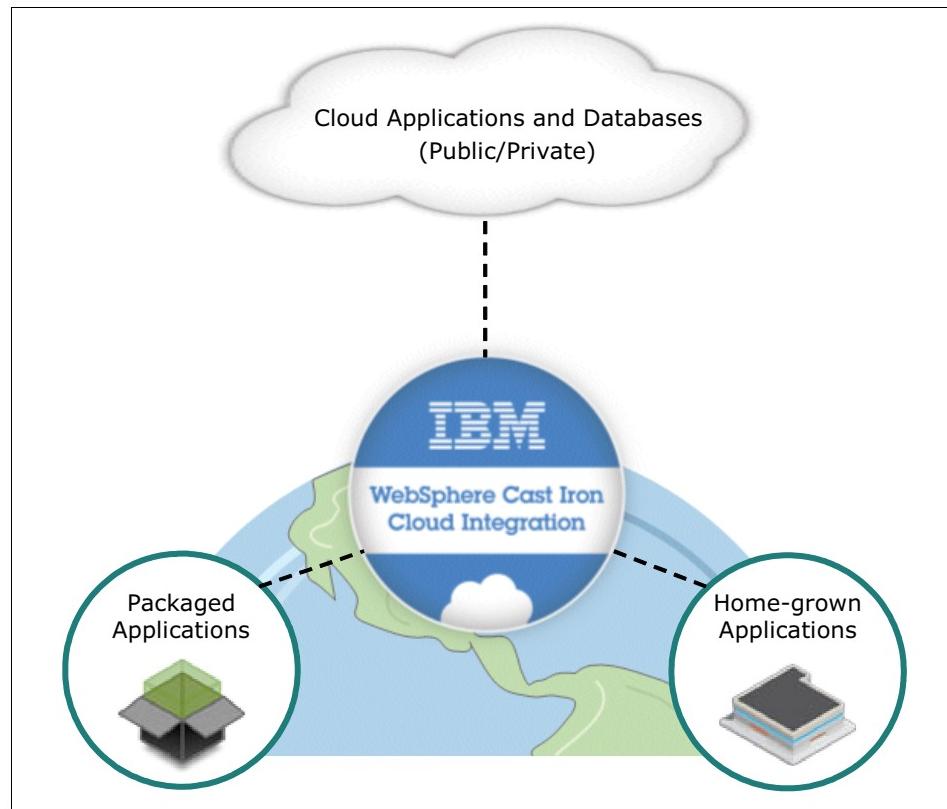


Figure 2-1 IBM WebSphere Cast Iron

The IBM WebSphere Cast Iron environment focuses on business, user, and application requirements to remove the daunting complexity of integration. Cast Iron uses a configuration approach instead of a coding approach to reduce the time needed to integrate systems on the cloud. This ensures there is a quick win for the business.

Cast Iron provides the following models of implementation:

- ▶ An appliance model where the appliance can be either physical hardware or a virtual machine and where the Integration Appliance is installed on premise, behind the firewall
- ▶ A SaaS model (Cast Iron Live)

Each of the three form factors (physical, virtual, or Cast Iron Live) allows for cloud-to-cloud, cloud-to-on-premise, and on-premise to on-premise integration. Each form factor allows for real time, near real time, and batch integrations.

The following section introduces the components of a Cast Iron implementation, differentiating between the on-premise Integration Appliance and the cloud model (Cast Iron Live).

2.1.1 Appliance model

An on-premise appliance normally resides behind the firewall and not within a DMZ. For the development, test, and production life cycles, the runtime environments are typically separated and each has its own Integration Appliance

that accesses the endpoints that are necessary for that environment. Figure 2-2 outlines the approach for using an on-premise Integration Appliance.

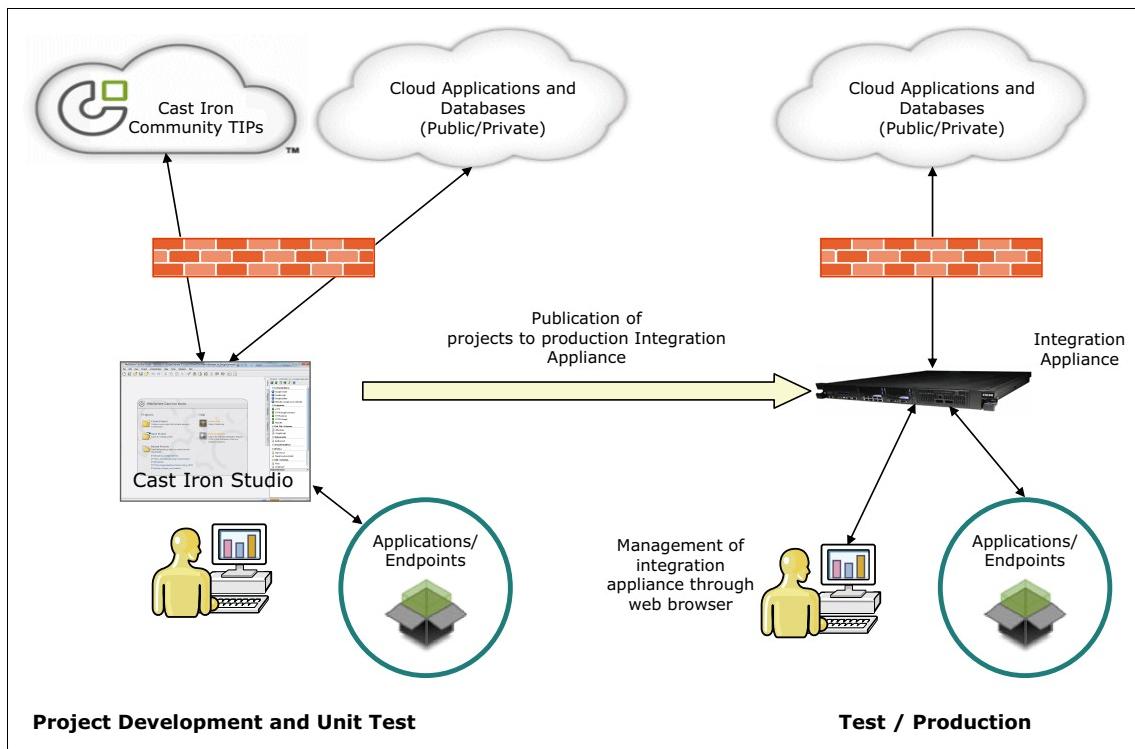


Figure 2-2 Architecture overview of the Cast Iron appliance model

Integration projects contain orchestrations that provide the logic of the integration. The integration projects are created and deployed using Cast Iron Studio. The Cast Iron Community maintains a solutions repository of Template Integration Projects (TIPs), that can be used as is or can be customized to orchestrate common integration types. After deployment, you have full control of projects and their orchestrations on the appliance through a Web Management Console.

The orchestrations are started through a mechanism that is defined by an activity in the orchestration. These activities include scheduling, polling, or waiting for an incoming request, such as an HTTP Receive Request.

The Integration Appliance requires connectivity to the endpoints that are required by the orchestration. Data flows through the Integration Appliance and is stored internally as XML variables. The Integration Appliance allows you to control the logging for each orchestration.

The Service Management Extensions for Hybrid Cloud will only work with either the physical or virtual options of the Integration appliance.

2.1.2 Cast Iron Live

The multi-tenant Cast Iron Live service includes the following key components that allow you to design, run, and manage integration projects, all in the cloud:

- ▶ A clustered runtime engine that runs the integration projects and that has built-in fault tolerance and recovery mechanisms
- ▶ A multi-tenant highly available system to store the designed integration projects
- ▶ A load balancer to intelligently manage the loads throughout the various runtime engines
- ▶ Highly available file systems to store and manage logs that are related to the integrations

The design environment, referred to in Cast Iron Live as *Designer*, provides the same design capabilities and user interface as Studio does for the on-premise

Integration Appliance. Figure 2-3 illustrates the architectural overview of the Cast Iron Live solution.

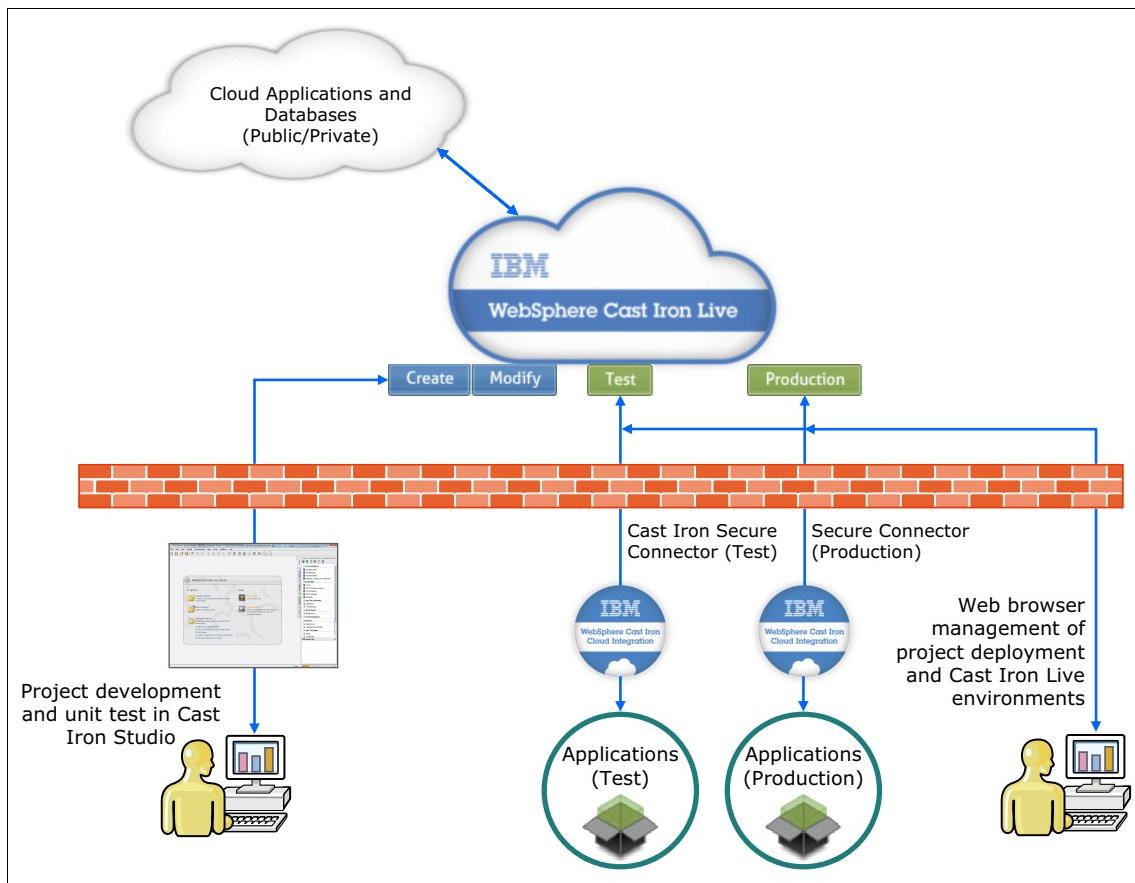


Figure 2-3 Overview of the architecture of the Cast Iron Live solution

You can access the Cast Iron Live interface from a web browser by typing in the following address <https://cloud2.castiron.com/login>

First-time users can trial, for up to 60 days, the Cast Iron Live model through the following website <https://express.castiron.com/express/#>

Attention: The Service Management Extensions for Hybrid Cloud will not work with the Cast Iron Live solution because it is a multi-tenant solution on the public cloud.

2.2 Examples of Cast Iron implementations

The following example use cases illustrate where WebSphere Cast Iron provides a solution for integration:

- ▶ Exposing ERP information to sales agents

A company needed to make sales order information that was stored in multiple ERP systems available to sales agents who used a SaaS-based CRM system. Integrating the ERP information with the CRM system provides the agents visibility to the necessary information and avoids the agents having to sign on to multiple systems. The company had only a small IT department with no substantial development capability, so they need a tool that business analysts can use.

The initial project involved using WebSphere Cast Iron to create three orchestrations that took less than two weeks to develop. The orchestrations provided real-time integration of the ERP information into the CRM system. Over time, the integration capabilities were expanded. There are now over 150 orchestrations connecting a multitude of systems that are running both in batch and on request.

- ▶ Integrating separate insurance policy systems

An insurance company has multiple systems that are used to manage new business. Having one system that provides quotes for policies and another system that is used to administer policies occasionally resulted in duplicate data entry and as a result, generated error conditions. In addition, many of the workflows involved in bringing in new business are manual, and vendors must provide documents in multiple formats.

The company used WebSphere Cast Iron to provide bidirectional integration between the two systems. Policy quote information is integrated into the administrative system and vice versa, making it necessary to enter data only one time and ensuring the integrity of the data in both systems. Integrating data through key business systems ensured that files were routed correctly and that workflow was not interrupted if a user is unavailable. Vendors must provide data only in one format, with the integration solution making the data available to the systems where it is needed in the proper format. Finally, the integration solution eliminates errors that were introduced because of duplicate data entry.

- ▶ Synchronizing e-commerce systems and linking to the SaaS CRM application to provide a single view of customers

A commercial enterprise had two e-commerce systems that are populated with customer and order data that is currently not synchronized. The enterprise also has a CRM application where customer data is kept, and a fulfillment system with inventory and product data. The enterprise used

WebSphere Cast Iron to integrate the data in these systems to provide a more cohesive view of customer data.

The initial implementation included the following orchestrations that provided the initial population of data throughout the systems and the synchronization of information between systems:

- Data providing a single view of the customer was populated into the SaaS CRM system from the enterprise's existing systems using a web services integration.
- An orchestration with a short polling interval pushes account, order, inventory, and product information to the single view from the e-commerce and fulfillment systems, ensuring that the customer service staff has the most current information.
- A nightly refresh of all stock keeping units in the e-commerce system occurs based on information from the fulfillment system. The orchestration uses secure file transfer to perform the refresh.
- Synchronization of customer records between the two e-commerce systems occurs on a regular basis, ensuring that when customers sign in to a system, their data is available to them.

2.3 How WebSphere Cast Iron integrates applications

WebSphere Cast Iron provides an approach to integrating applications that does not require any programming knowledge. You can build integration flows in WebSphere Cast Iron Studio, which is a graphical development environment that is installed to a personal computer (PC).

With WebSphere Cast Iron Studio, you create an integration project that contains one or more orchestrations. Each orchestration is built with a number of activities

that define the flow of data. You can define the details of an activity from the configuration panes within WebSphere Cast Iron, as shown in Figure 2-4.

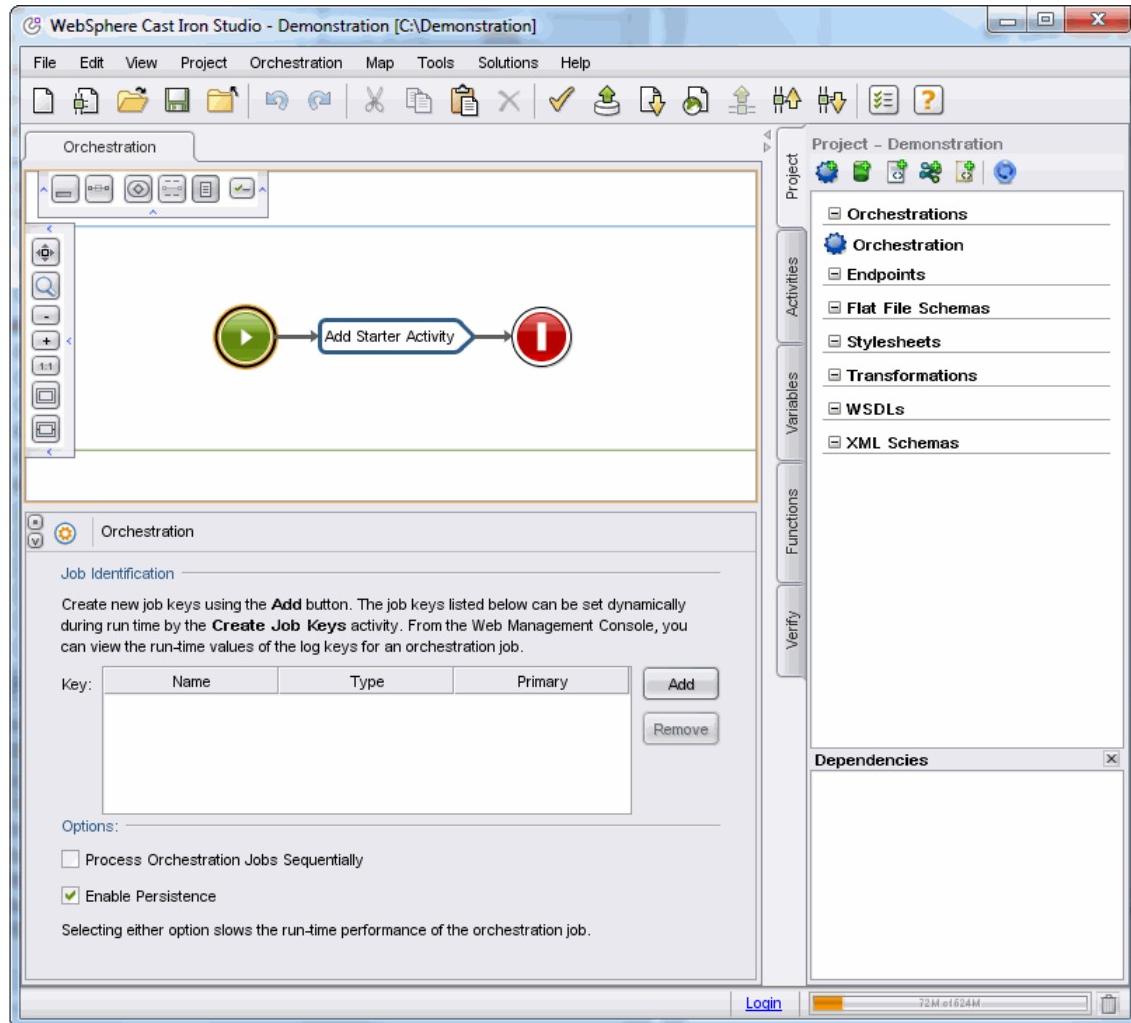


Figure 2-4 Cast Iron Studio

A project contains all of the assets that are required for the orchestrations to run, including any file schemas, WSDL files, and functions. The project also defines connectivity to the sources of data, the *endpoints*. Cast Iron contains many built-in connectors to applications (for example, SAP), databases, and web services that make connecting to these endpoints straightforward.

Cast Iron Studio provides easy-to-use tools for configuring properties and for mapping input data to output data. Cast Iron also provides Template Integration Projects (TIPs) that encapsulate a specific integration use case between specific endpoints and that include preferred practices. You can download these TIPs from the Cast Iron community and modify to fit your precise needs. Figure 2-5 outlines the Cast Iron approach to creating orchestrations.

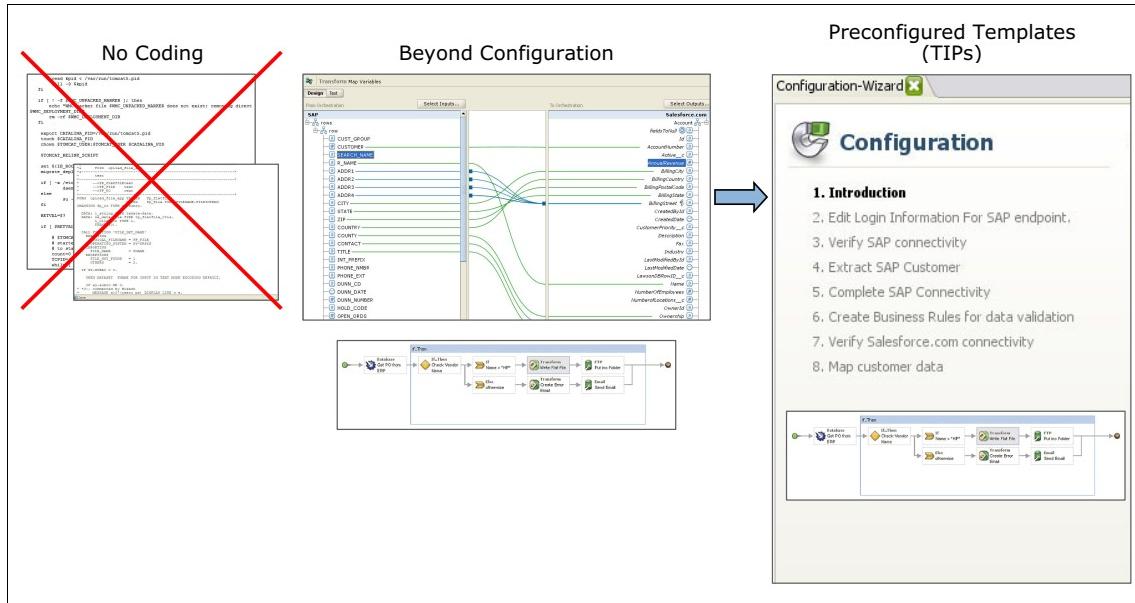


Figure 2-5 Cast Iron approach

Within Studio, there is a simulated runtime environment to enable unit testing of the orchestrations with visibility of all data at all points. After you test the project, you can publish the project to the run time, where you can then manage and monitor the project. A browser-based interface, the Web Management Console (WMC), provides the capability to manage all functions on the run time.

As with any tool, there are preferred practices to implement designs. When processing large amounts of data, a poor design can mean that the integration takes an extraordinary amount of time or that it simply does not work.

Each endpoint is different in its requirements, and any integration flow must use the characteristics of the endpoint to its best advantage, for example processing batches of data to optimize the speed of transactions.

2.4 Service Management Extensions for Hybrid Cloud

The Service Management Extensions for Hybrid Cloud extends the on-premises suite of products through integration of public and private clouds, and intra-cloud environments that involve multiple distributed data centers and distributed workloads with the following functionalities:

- ▶ Monitor cloud resources and on-premises resources in the same user interface
- ▶ Provision cloud resources and on-premises resources in the same user interface
- ▶ Govern where resources are used, on-premises or in the cloud
- ▶ Ensure proper authorization to use services
- ▶ Synchronize user directories so that on-premises users can access the cloud

The Service Management Extensions for Hybrid Cloud delivers these functionalities on either the IBM WebSphere DataPower® Cast Iron Appliance XH40, which is a physical appliance, or the IBM WebSphere Cast Iron Hypervisor Edition, which is a virtual appliance. It does not run on the Cast Iron Live or Express editions. The physical and virtual Cast Iron Appliance is referred to as the “WebSphere Cast Iron Appliance.”

2.4.1 Software components of the solution

The software components of the Service Management Extensions for Hybrid Cloud solution are listed here:

- ▶ The Provisioning Plug-in enables provisioning of resources into IBM SmartCloud™ Enterprise and Amazon Elastic Compute Cloud.
- ▶ The Monitoring Plug-in enables monitoring of resources in the IBM SmartCloud Enterprise and Amazon Elastic Compute Cloud.
- ▶ The Directory Synchronization Plug-in enables synchronizing enterprise LDAP user information with LotusLive user information.

2.4.2 Prerequisites

Table 2-1 on page 26 provides a breakdown of requirements for each usage scenario of the Service Management Extensions for Hybrid Cloud.

Table 2-1 Prerequisite products for each scenario

| Product | Version | Provision resources | Monitor resources | Govern and manage workload | Synchronize directories |
|--|-----------------|---------------------|-------------------|----------------------------|-------------------------|
| WebSphere DataPower Cast Iron XH40 Appliance | 6.1.0.3 | Y | Y | Y | Y |
| WebSphere Cast Iron Hypervisor Edition | 6.1.0.3 | Y | Y | Y | Y |
| IBM LotusLive | | | | | Y |
| Tivoli Monitoring | 6.2.2 | | Y | Y | |
| Tivoli Monitoring Agent | 6.2.2 FixPack 2 | | Y | Y | |
| Tivoli NetCool/Impact | 5.1.1 FixPack 1 | | | Y | |
| Tivoli Service Automation Manager | 7.2.1 FixPack 4 | Y | | Y | |
| IBM WebSphere ILOG JRules | 7.1.1 | | | Y | |

2.5 Usage scenarios

Some of the most common scenarios where this hybrid solution is applicable are listed here and explained in the following sections:

- ▶ New client
- ▶ Clients with an existing WebSphere Cast Iron installation
- ▶ Clients with Tivoli Monitoring
- ▶ Clients with Tivoli Service Automation Manager
- ▶ Clients with LDAP

2.5.1 New client scenario

In the “new client” scenario, a customer wants to extend its network to leverage the provisioning, monitoring, or both, of public clouds. An IBM WebSphere DataPower Cast Iron Appliance XH40 or WebSphere Cast Iron Hypervisor Edition is not deployed and Service Management Extensions for Hybrid Cloud prerequisite software is not installed.

The solution consists of the following tasks:

1. Install and configure the Tivoli Service Delivery Manager software.
2. Deploy the IBM WebSphere DataPower Cast Iron Appliance XH40 or WebSphere Cast Iron Hypervisor Edition.
3. Install the Service Management Extensions for Hybrid Cloud product software.
4. Contact Amazon Web Services (AWS) and or IBM Smart Cloud Enterprise services to request a user account.
5. Configure each hybrid service. Part 2, “Usage scenarios” on page 59 explains the details of each usage scenario.

2.6 Clients with a WebSphere Cast Iron setup scenario

In this scenario, a client has deployed IBM WebSphere DataPower Cast Iron Appliance XH40 or WebSphere Cast Iron Hypervisor Edition, but Service Management Extensions for Hybrid Cloud prerequisite software is not installed. The customer wants to provision and monitor resources on public clouds.

The solution consists of the following tasks:

1. Upgrade to Cast Iron firmware version 6.1.0.3.
2. Install and configure the Tivoli Service Delivery Manager software.
3. Install the Service Management Extensions for Hybrid Cloud product software.
4. Contact Amazon Web Services (AWS) and or IBM Smart Cloud Enterprise services to request a user account.
5. Configure each hybrid service.

2.6.1 Clients with a Tivoli Monitoring setup scenario

In this scenario, a customer has deployed Tivoli Monitoring and wants to extend their Tivoli Monitoring and add provisioning capabilities into public clouds.

The solution consists of the following tasks:

1. Upgrade to Tivoli Monitoring version 6.2.2 or later.
2. Install and configure the Tivoli Service Automation Manager.
3. Install the Service Management Extensions for Hybrid Cloud product software.
4. Contact Amazon Web Services (AWS) and or IBM Smart Cloud Enterprise services to request a user account.
5. Configure each hybrid service.

2.7 Clients with a Tivoli Service Automation Manager setup scenario

In this scenario, a customer has deployed Tivoli Service Automation Manager and wants to leverage all of Service Management Extensions for Hybrid Cloud capabilities.

The solution consists of the following tasks:

1. If the current version of Tivoli Service Automation Manager is Version 7.2.0 or earlier, upgrade to Version 7.2.1.
2. Install and configure Tivoli Monitoring, WebSphere ILOG JRules, and Tivoli Netcool/Impact.
3. Deploy IBM WebSphere DataPower Cast Iron Appliance XH40 or WebSphere Cast Iron Hypervisor Edition.
4. Install the Service Management Extensions for Hybrid Cloud product software.
5. Install the monitoring, provisioning, and LotusLive Service Management Extensions for Hybrid Cloud plug-ins.
6. Contact Amazon Web Services (AWS) and or IBM Smart Cloud Enterprise services and Lotus Live services to request a user account.
7. Configure each hybrid service.

2.7.1 Clients with a LDAP v3-compliant Directory Server scenario

In this scenario, a customer has deployed an LDAP v3 compliant directory server and wants to manage LotusLive accounts.

The solution consists of the following tasks:

1. Deploy IBM WebSphere DataPower Cast Iron Appliance XH40 or WebSphere Cast Iron Hypervisor Edition.

The IBM WebSphere DataPower Cast Iron Appliance XH40 or WebSphere Cast Iron Hypervisor Edition includes the required version of Tivoli Directory Integrator preinstalled.

2. Contact LotusLive services to request a user account.
3. Install and configure the LotusLive Service Management Extensions for Hybrid Cloud plug-ins, which also installs the required Tivoli Directory Integrator run-time engine.

A Tivoli Directory Integrator Server in a customer site, if available, is not affected.

2.8 Installing the Service Management Extensions for Hybrid Cloud

The prerequisite tasks to be completed prior to installing the Service Management Extensions for Hybrid Cloud are listed here:

- ▶ Create a public cloud account at one of the following public cloud providers.
 - IBM Smart Cloud Enterprise
 - Amazon Elastic Compute Cloud (Amazon EC2)
 - Lotus Live (only for Directory Synchronization)
- ▶ Set up the WebSphere Cast Iron system (only the physical or virtual appliances are compatible).
- ▶ Install the prerequisite software based on your usage scenario.
 - Tivoli Service Automation Manager
 - Tivoli Monitoring
 - Tivoli NetCool/Impact
 - WebSphere ILOG JRules

The next step is to download the Extensions from the following website:

<https://www-304.ibm.com/software/brandcatalog/ismlibrary/details?catalog.label=1TW10TS0D>

2.8.1 Installing the extensions through the Cast Iron design console

This process assumes that you have a virtual appliance up and running.

1. Log on to the WebSphere Cast Iron Design Console at <http://hostname/express> (replace the histamine with the IP address of your device) with the admin user account and relevant password (Figure 2-6).



Figure 2-6 WebSphere Cast Iron Design Console logon through web browser

2. After you successfully log in to the Design Console, install the plug-ins from the extracted files that were originally downloaded from the Service Management Extensions site, as shown in Figure 2-7.

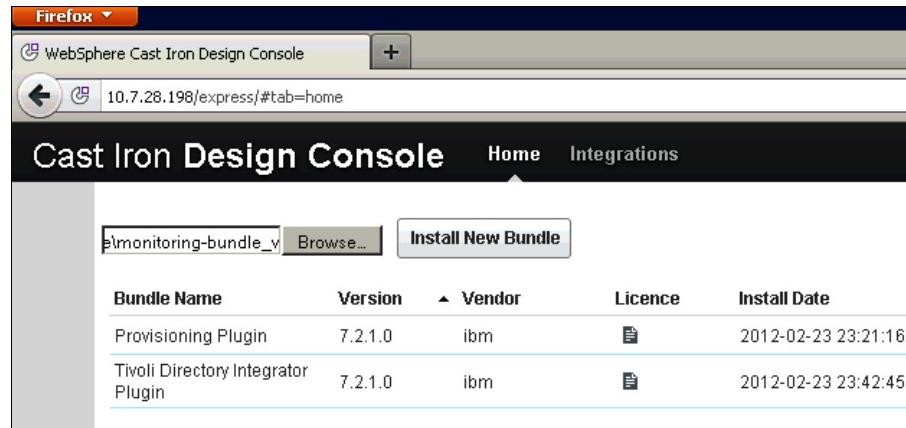


Figure 2-7 Click the “Install New Bundle” button to start the installation

After the extensions are successfully installed, a window similar to that shown in Figure 2-8 is displayed.

The screenshot shows the same interface as Figure 2-7, but now with four entries in the table, indicating successful installation of the third extension.

| Bundle Name | Version | Vendor | Licence | Install Date |
|------------------------------------|---------|--------|---------|---------------------|
| Provisioning Plugin | 7.2.1.0 | ibm | ■ | 2012-02-23 23:21:16 |
| Tivoli Directory Integrator Plugin | 7.2.1.0 | ibm | ■ | 2012-02-23 23:42:45 |
| Monitoring | 7.2.1.0 | ibm | ■ | 2012-02-23 23:23:20 |

Figure 2-8 All three extensions have been installed and running successfully

3. Create the integration between the respective plug-ins and your service environment, as shown in Figure 2-9.

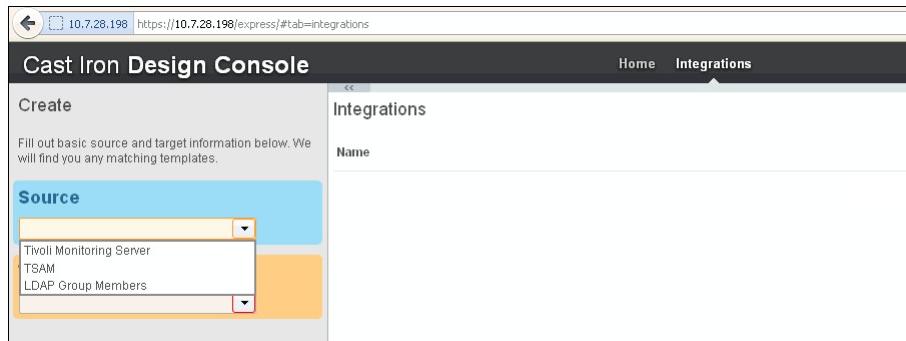


Figure 2-9 Creating the integration between the plug-ins and your service environment

Creating the integration between the source and target through the plug-ins is discussed in more detail in Part 2, “Usage scenarios” on page 59.

2.9 Example scenario

This section provides a high-level look at an example scenario illustrating where and when this solution can be used. This scenario is explored in more detail in Part 2, “Usage scenarios” on page 59.

The scenario is based on a fictional, large transnational retail company called IBM Redbooks Store with 1000 stores. This company acquires another retail company with 100 stores to add to their portfolio of products.

2.9.1 Provisioning

After the acquisition, the company realized it needed an integrated inventory management system so that all 1100 stores could sell the complete product line. To ensure that the new line of products is sold at all 1100 stores, the company had an immediate need to set up an integration and test environment to support the new integrated inventory management system.

The company leveraged a Hybrid Cloud model to obtain the additional capacity needed for the brief period of time during which they will work on integration and test, as they expanded their capacity of their on-premise cloud resources to host the production workload.

As illustrated in Figure 2-10, the company will be able to utilize its existing Tivoli Service Automation Manager system to create the new testing environment on the cloud through the Provisioning Bundle installed on its WebSphere Cast Iron appliance.

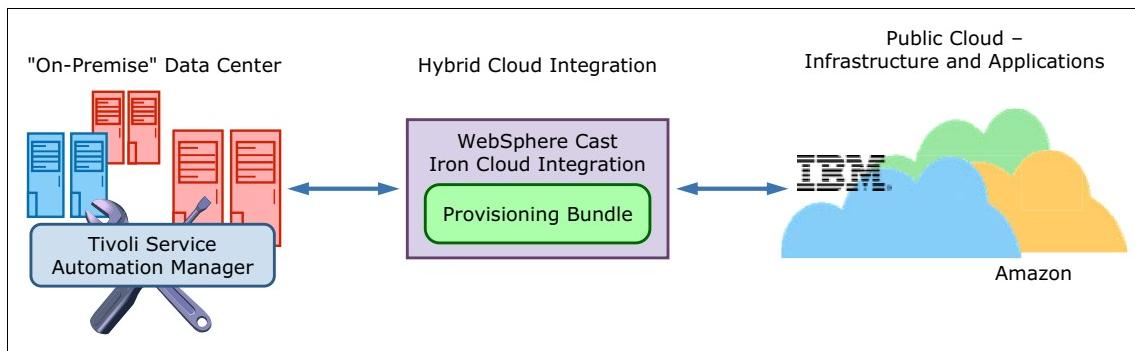


Figure 2-10 Overview of the Provisioning integration

The end result is an enabled federated service catalog as shown in Figure 2-11 on page 34. This federated catalog will help the administrator to create, manage

and tear down server resources on the cloud using the same Tivoli Service Automation Manager that is currently used internally.

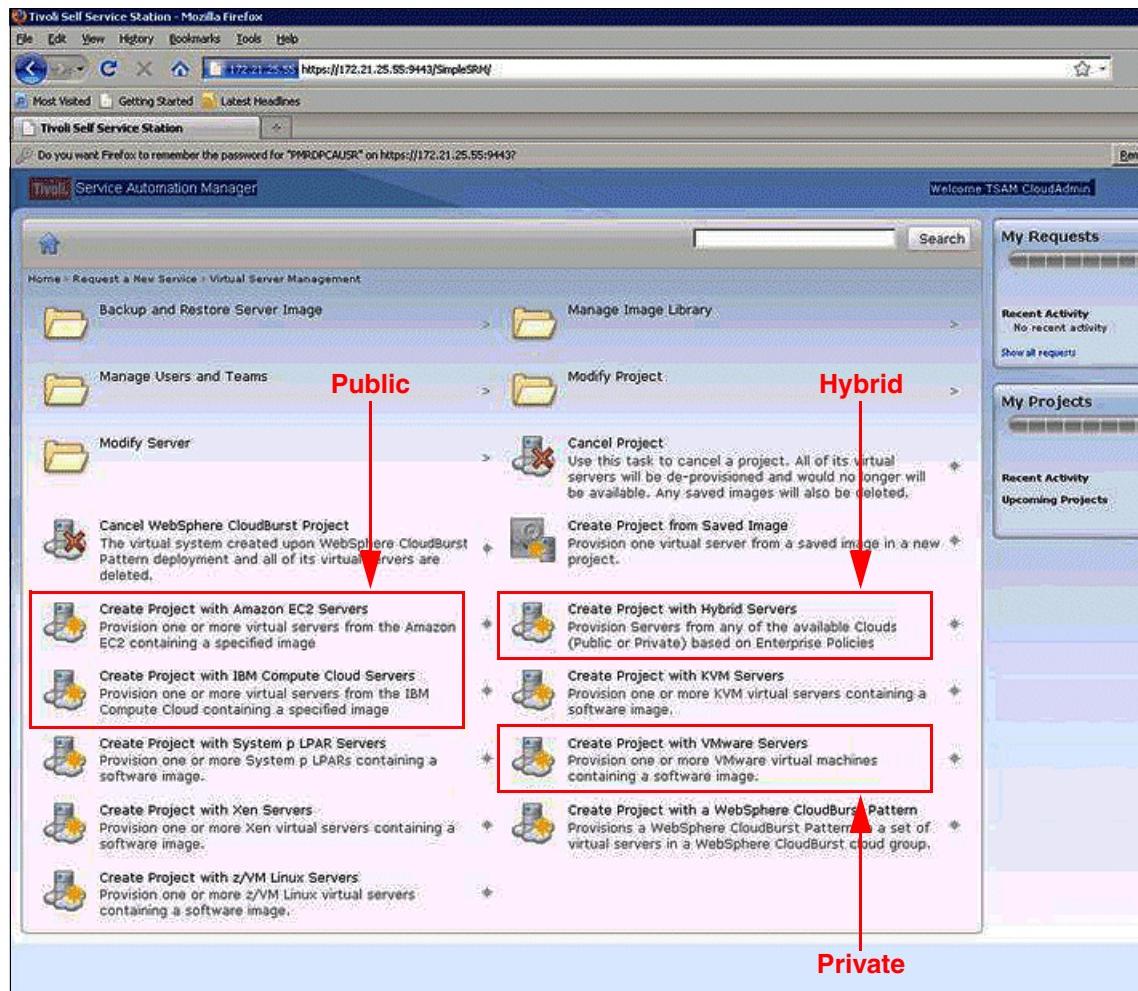


Figure 2-11 Federated service catalog

Chapter 5, “Provisioning scenario” on page 93, provides detailed information about this solution, and explains the installation and configuration steps.

2.9.2 Monitoring

During performance testing on its new inventory management system, the company became aware that the new product line is a seasonal product. That is, although the product line is sold throughout the year, it sells most heavily during October, November, and December.

When the company ran its performance tests with simulated data for a year, it realized that it will need additional CPU capacity in October, November, and December.

Due to this cyclical period of overload, the company had to utilize the public cloud providers. As a result, it used its Tivoli Monitoring system to monitor both its on-premise data center and its off-premise systems residing on the cloud, as shown in Figure 2-12.

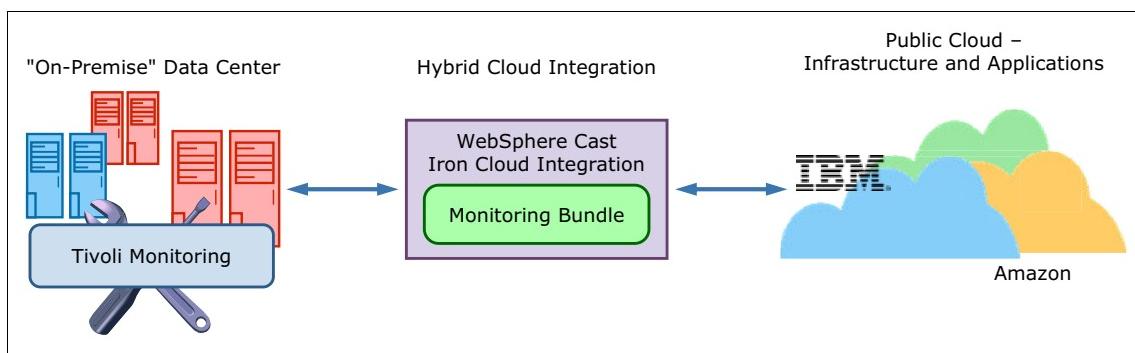


Figure 2-12 Overview of the Monitoring integration

Figure 2-13 shows the end result when viewed through Tivoli Monitoring. The new unified workload monitoring capability ensures that the IT team will be able to monitor both data centers through one interface, and it enables them to increase the system resources on demand.

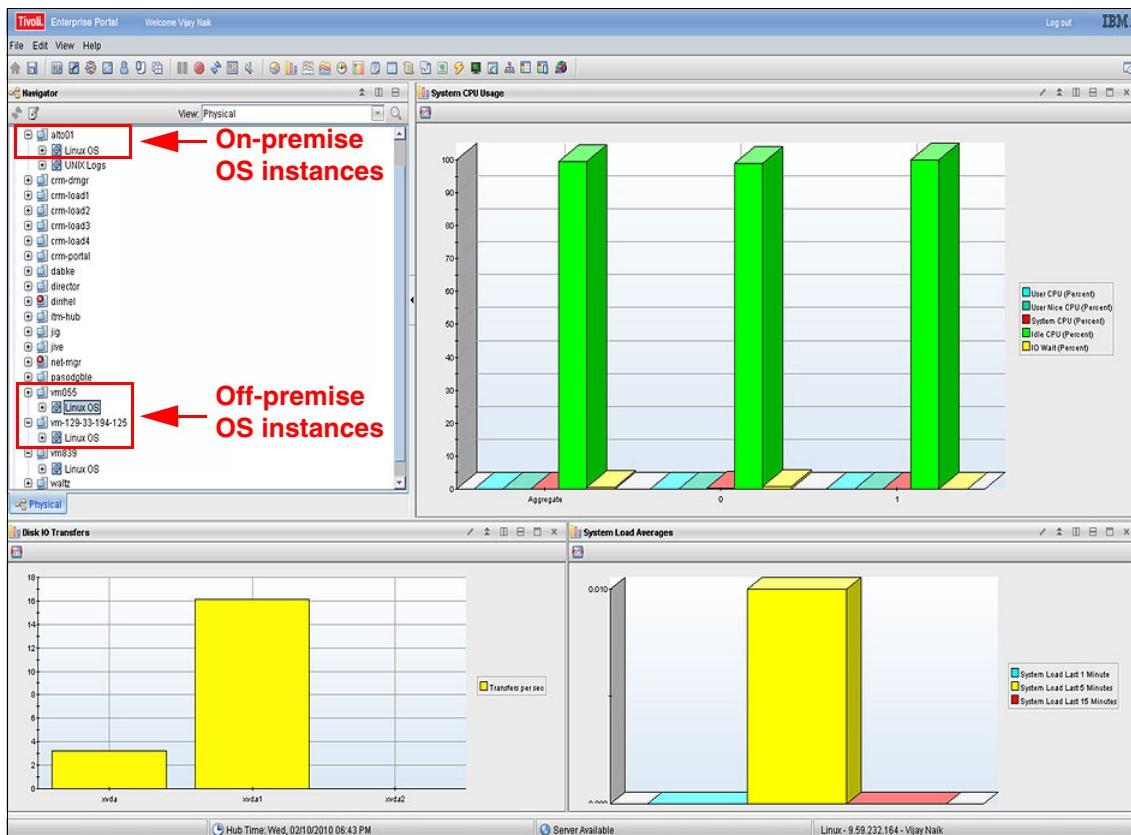


Figure 2-13 Unified Workload Monitoring through Tivoli Monitoring and Cast Iron

Chapter 4, “Monitoring scenario” on page 61, explains the installation and configuration steps involved in getting this system up and running.

2.9.3 Directory Synchronization

In addition to the larger inventory in October, November and December, the company hires additional staff during these months. This requires the provisioning of additional users in the company's internal LDAP directory.

IBM Redbooks Store uses Lotus Live (<https://www.lotuslive.com/en/>) as the email system for the temporary staff every year, and it needs to synchronize its internal LDAP directory with the temporary staff users in LotusLive.

This is achieved using the Directory Synchronization bundle installed on the company's existing Cast Iron appliance, and integrating it to the Lotus Live system on the cloud as shown in Figure 2-14.

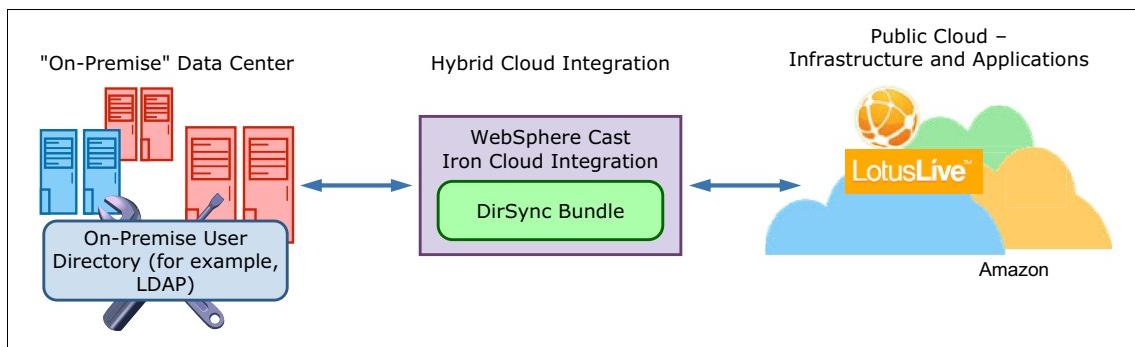


Figure 2-14 Overview of the Directory Synchronization integration

Chapter 6, “Directory Synchronization scenario” on page 153, explains the steps needed to get the solution up and running.

2.9.4 Workload Governance and Management

When IBM Redbooks Store implements the new inventory management system that supports all 1100 stores, the company will use a Hybrid Cloud that supports the entire workload during most quarters with on-premise resources, and then leverage public resources during the heaviest months.

In addition, the company will save additional money by decommissioning users and compute resources after the holiday business season ends and business returns to normal volumes.

Figure 2-15 gives an overview of the policy-based system that enables the business to create system resources based on business rules.

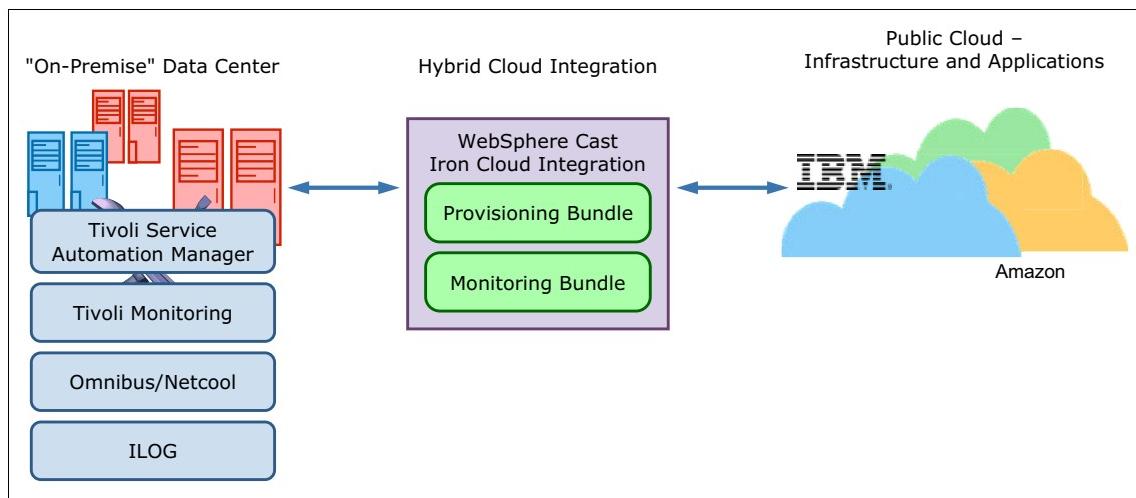


Figure 2-15 Workload Governance and Management overview

This solution provides the business with the following capabilities:

- ▶ To create, manage, and tear down server resources based on business policies
- ▶ To optimize resource allocation based on workload

The WebSphere ILOG Rules engine is used to define the workload placement, as shown in Figure 2-16.

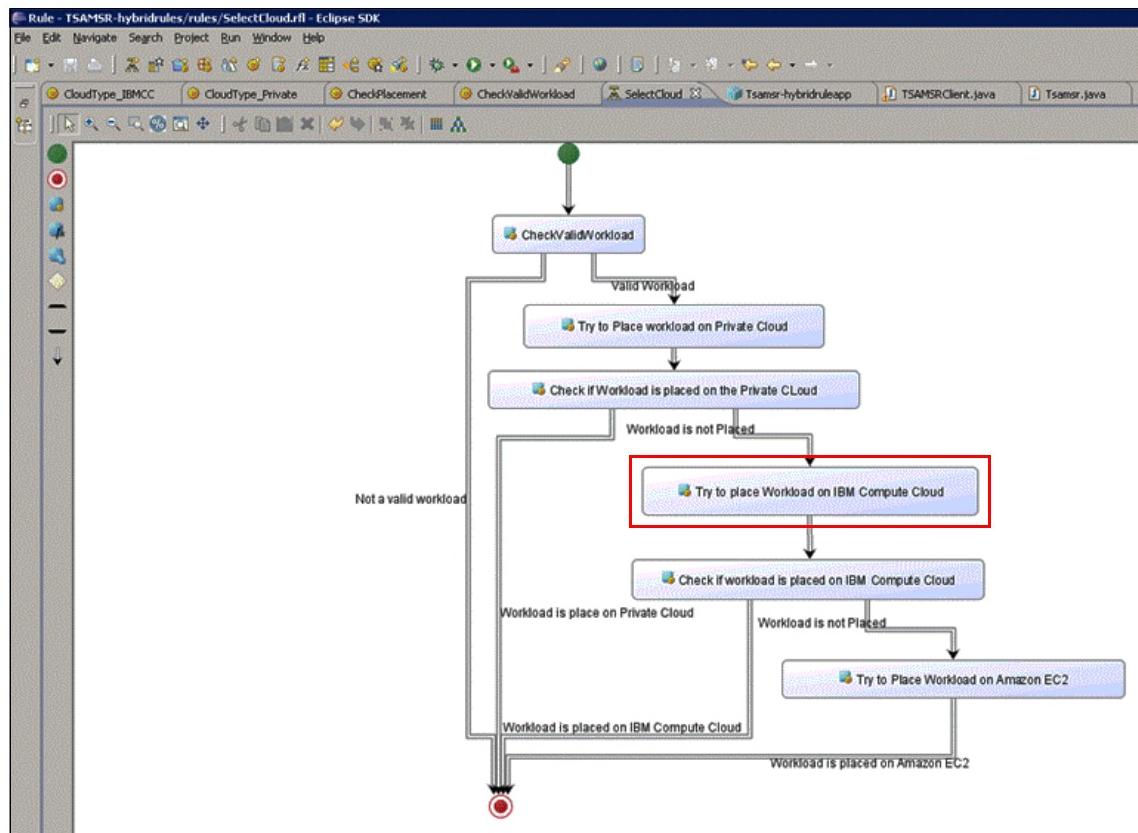


Figure 2-16 Using the WebSphere ILOG Rules engine for workload placement

The rule flow determines which cloud to place the workload on. For example, in Figure 2-16 the “Try to place Workload on IBM Compute Cloud” task determines whether the workload should be placed on the IBM cloud with the following rule:

PMRDPCLOUD WORKLOADTYPE of the Service Request contains “Linux OS for TEST”

THEN

set the PMRDPCLOUD PLACEMENT of the Service Request to IBM SCE’;

Chapter 7, “Workload Governance and Management scenario” on page 173, explains the steps needed to get this solution up and running.

2.10 Summary

The IBM Service Management Extensions for Hybrid Cloud provides the following features:

- ▶ Monitoring Integration with IBM Tivoli Manager (ITM)
An integrated monitoring dashboard creates a single view for monitoring on-premise and off-premise resources.
- ▶ Provisioning and governance integration with Tivoli Service Automation Manager (TSAM)
This feature provides the ability to request and provision IBM Smart Cloud Enterprise and Amazon hybrid resources through TSAM's graphical user interface.
- ▶ User Directory Integration for Lotus Live with Tivoli Directory Integrator
Automated user account management is possible between enterprise and Lotus Live.
- ▶ Workload Governance and Management
You can set resource overflow and underutilization thresholds on Event Correlation Service on Impact, TSAM rules on ILOG Business Rules and TSAM escalations and actions.
Workload resources can be automatically balanced based on the dynamics of the system load.



IBM Cloud Service Management Platform

In this chapter we provide an overview of the IBM Cloud Service Management Platform (CCMP), which provides both business and operational support services to service consumer and development portals. This common Cloud Service Management Platform enables both cloud service consumers and cloud service creators to use a robust platform to consume or create cloud services. In addition, we discuss how the roles of cloud service consumer, cloud service provider, and cloud service creator interact and benefit from a common cloud service management platform.

Next, we discuss how the Cloud Service Management Platform applies and supports hybrid cloud environments. Finally, we cover how the Service Management Extensions help support and extend the Cloud Service Management Platform.

In this chapter we discuss the following topics:

- ▶ Cloud computing architecture
- ▶ Mapping IBM products to the CCMP
- ▶ Solutions supporting the Business Support Services (BSS)
- ▶ Solutions supporting the Operational Support Services (OSS)
- ▶ Cloud Service Management Platform for Hybrid Cloud

3.1 Cloud computing architecture

The IBM Cloud Reference Architecture (Figure 3-1) is designed to provide guidance for cloud computing and provide a blueprint for effective security, resiliency, service management, governance, business planning and lifecycle management. These components of the cloud reference architecture enable an enterprise to control the cloud environment most effectively, optimize productivity, reduce associated labor costs and ensure a safe environment for business users.

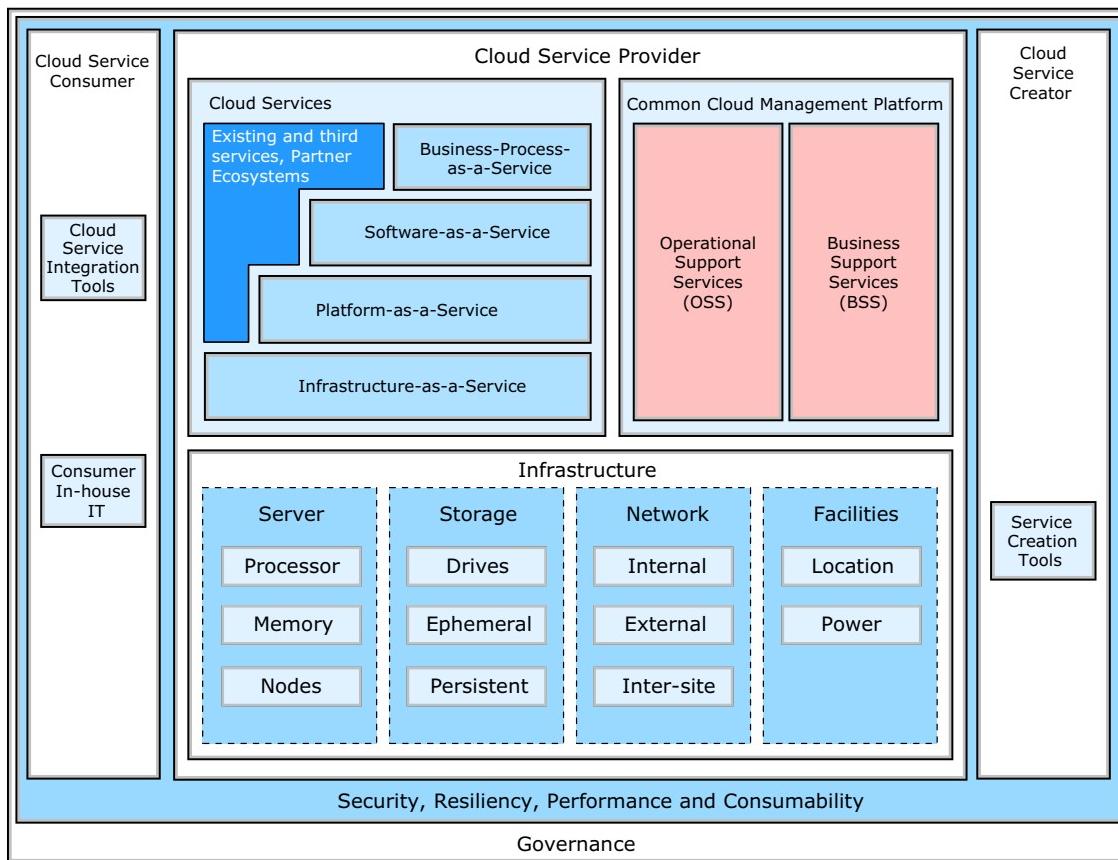


Figure 3-1 Cloud Reference Architecture

3.1.1 Key roles in a cloud ecosystem

There are three key roles that support a cloud ecosystem.

First, the *cloud service consumer* role represents an organization, a human being, or an IT system that consumes service instances delivered by a particular cloud service. The service consumer might be billed for its interactions with cloud services and the provisioned service instances. A service consumer can also be viewed as a kind of super-role representing the party consuming services. The cloud service consumer browses the service offering catalog and triggers the provisioning and management of cloud services.

Next, the *cloud service provider* role has the responsibility of providing cloud services to cloud service consumers. A cloud service provider is defined by the ownership of a Common Cloud Management Platform (CCMP). This ownership can either be realized by truly running a CCMP by the provider itself, or consuming one as a service. A person acting in the role of a cloud service provider and a cloud service consumer at the same time is a partner of another cloud service provider who is reselling cloud services or consuming cloud services and adding value add functionality on top, which in turn is provided as a cloud service.

Finally, the *cloud service creator* role is responsible for creating a cloud service, which can be run by a cloud service provider and exposed to cloud service consumers. Typically, cloud service creators build their cloud services by leveraging functionality which is exposed by a cloud service provider.

Management functionality, which is commonly needed by cloud service creators, is defined by the CCMP architecture. A cloud service creator designs, implements, and maintains runtime and management artifacts specific to a cloud service. Just as with the cloud service consumer and the cloud service provider, the cloud service creator can be an organization or a human being.

Typically, the operations staff is responsible for operating a cloud service and will be closely integrated with the development organization that is developing the service.

3.1.2 Common Cloud Management Platform

The Common Cloud Management Platform (CCMP) leverages four key design principles to provide a complete service management solution for cloud services.

First, the CCMP supports the ability to design cloud services for scale. This is known as the “Efficiency Principle” and has the primary objective of driving down costs and time to response by orders of magnitude by leveraging cloud characteristics such as elasticity, self-service access, and flexible sourcing.

The second principle is known as the “Lightweight Principle” and supports lean and lightweight service management policies, processes, and procedures. This

principle allows for standardization in cloud environments to reduce management costs.

The next design principle that the CCMP leverages is the “Economies-of-Scale” principle, which maximizes the sharing of management, infrastructure, and platform components across cloud services. This sharing of components allows for a reduction in costs and time to market.

Finally, the CCMP leverages the “Genericity Principle” which allows for service templates to manage service instances generically along their lifecycle. This generic design of components allows for a single management platform.

These four key design principles are leveraged to create the Operational and Business support services that make up the Common Cloud Management Platform.

Operational Support Services

The Operational Support Services (OSS) module (Figure 3-2) defines the set of systems management services that may be used by cloud service developers.

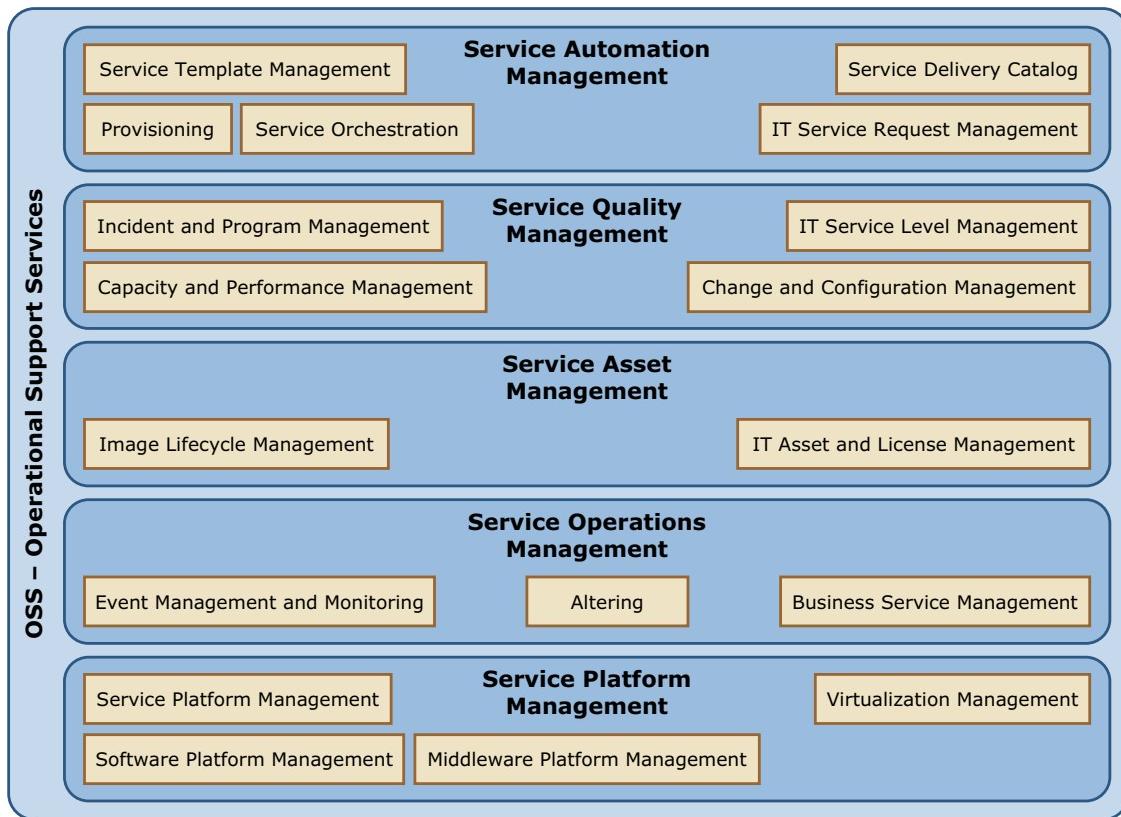


Figure 3-2 OSS architecture

Many of the management domains specified in the OSS can also be encountered in traditionally managed data centers such as monitoring and event management, provisioning, and incident and problem management. Although these management domains are conceptually the same in both traditional and cloud IT environments, in a cloud architecture these domains can be implemented in different ways.

Take, for example, incident and problem management. If a physical server fails in a traditional IT environment, a trouble ticket is opened and assigned to a systems administrator for manual intervention and resolution. If the problem is not resolved in the time specified by service level agreements (SLAs), it may be escalated until it is resolved.

However, if a physical server fails in a cloud environment, a new virtual machine or application is automatically brought up on another physical server where resources are available, without delay. This approach is typically referred to as “replace versus repair,” based on the assumption that manually repairing things is often more complex than replacing them. It is this kind of rules-based response to hardware failure, which allows operation at a much lower cost, that characterizes service management in the cloud.

Business Support Services

The Business Support Services (BSS) module (Figure 3-3) defines the capabilities required to enable the business management of one or more specific managed cloud services.

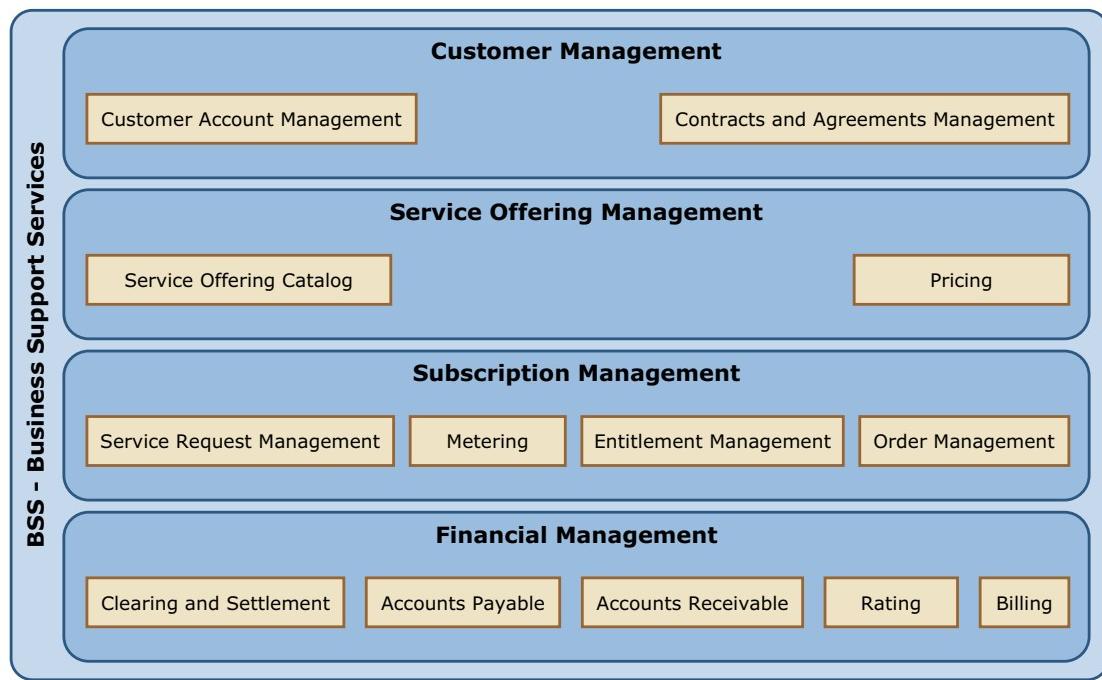


Figure 3-3 BSS architecture

Cloud services represent any type of IT capability that is offered by the cloud service provider to cloud service consumers. Typical categories of cloud services are infrastructure, platform, software or business process services. In contrast to traditional IT services, cloud services have attributes associated with cloud computing, such as a pay-per-use model, self-service acquisition of services, flexible scaling, and sharing of underlying IT resources. The CCMP is designed to enable these cloud-specific capabilities through the OSS, while handling the

ongoing management of all of the provider's cloud service instances. The BSS is responsible for handling all business-relevant aspects of a cloud service.

For example, the billing service component of the BSS must be capable of performing billing for the consumption of virtual machine resources, a multi-tenancy capable middleware platform (platform as a service) and a multi-tenancy application such as collaboration or customer relationship management (software as a service).

Other components of the BSS address service management and automation at the user interface level. These capabilities include a user-friendly self-service interface and service offering catalog through which consumers select, configure, arrange payment for and discontinue cloud services.

The BSS delivers the functions required to operate a self-service cloud business. These include automated, rules-based execution of pricing, contracts and agreements, invoicing, clearing and settlement. The BSS also provides business management capabilities, including offering, customer, subscriber, order, fulfillment and entitlement management. This enables the business side of the cloud services paradigm.

3.2 Mapping IBM products to the CCMP

The IBM Cloud Service Management Platform provides business and operational support services for the CCMP. In this section we discuss how IBM products map to the components in the CCMP, including:

- ▶ Solutions supporting the Operational Support Services
- ▶ Solutions supporting the Business Support Services

The product mapping is shown in Figure 3-4.

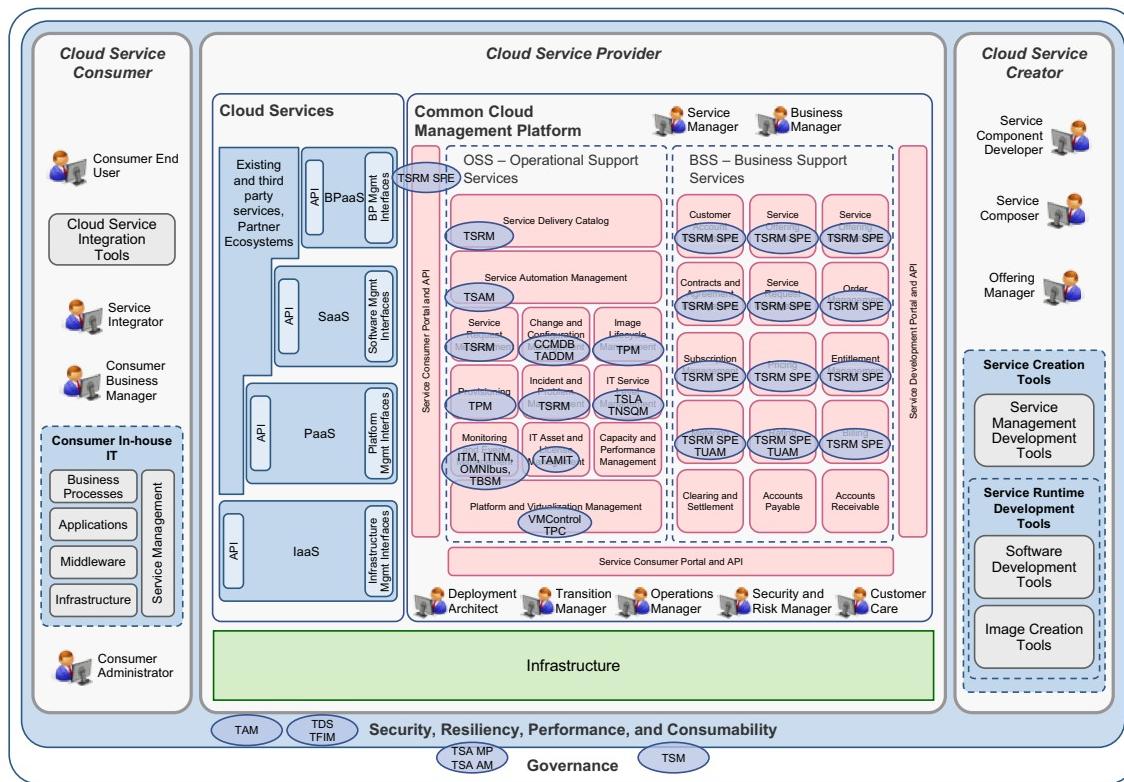


Figure 3-4 Cloud architecture-to-solution mapping

3.2.1 Solutions supporting the BSS

There are two Tivoli products that support the BSS, Tivoli Service Request Manager® for Service Providers and Tivoli Usage and Accounting. Together these products provide the BSS functions as defined in the Cloud Service Management Platform.

Tivoli Service Request Manager for Service Providers

Tivoli Service Request Manager for Service Providers helps service providers provide service delivery capabilities for multiple customers in a single deployed instance. This product includes the following key features:

- ▶ Reduces TCO by leveraging a single deployed instance to manage multiple customers

- ▶ Improves efficiency of service delivery with automatic notification and automatic assignments
- ▶ Manages multiple customers in many physical locations and provide unique customer agreements and rules to define entitlement of services and pricing
- ▶ Provides detailed and accurate billing with a review and approval cycle

For more information visit:

<http://www-01.ibm.com/software/tivoli/products/service-request-manager-sp/index.html>

IBM SmartCloud Cost Management

IBM SmartCloud Cost Management, previously known as Tivoli Usage and Account Manager, accurately assesses shared computing resource usage. This allows organizations to understand their costs by tracking, allocating, and invoicing by department or user.

This product includes the following key features:

- ▶ Collect, analyze, and bill based on usage and costs of shared Microsoft Windows, UNIX, Linux, IBM i5/OS™ and VMware computing resources
- ▶ Deliver detailed information and reports about the intricate use of shared resources while masking the underlying complexity
- ▶ Transform raw IT data into business information for cost allocation that spans business units, cost centers, applications, and users
- ▶ Consolidate a wide variety of usage data with Data Collectors and a powerful “business rules-driven” capability
- ▶ Automate daily cloud and non-cloud collection and billing operations for easy administration

For more information visit:

<http://www-01.ibm.com/software/tivoli/products/usage-accounting/>

3.2.2 Solutions supporting the OSS

The CCMP is supported by multiple Tivoli products, each providing a different solution for functions in the OSS.

- ▶ Change and Configuration Management Database and Tivoli Application Dependency Discovery Manager
- ▶ Tivoli Provisioning Manager

- ▶ Tivoli Service Level Advisor and Tivoli Netcool Service Quality Manager
- ▶ Tivoli Asset manager for IT
- ▶ VMControl and Tivoli Productivity Center
- ▶ Tivoli Service Automation Manager

Tivoli Change and Configuration Management Database

Tivoli Change and Configuration Management Database provides automated impact analysis and enables the quickest access to critical data that might cause outages due to change. This product includes the following key features:

- ▶ Workflow-driven change processes, which reduce costs by making it quick to process a change, and reduce risk by planning and communicating the change before it actually occurs.
- ▶ Blackout periods identify critical business periods when outages are expensive. Blackout periods can be restricted or locked down for maximum flexibility.
- ▶ Multi-customer enablement or Service Provider support change and configuration management for multiple customers in a single deployment.
- ▶ Configuration management database accuracy, which is ensured through complete audit and remediation processes that are ready for immediate use.
- ▶ Visualization tooling, which allows for topology viewing, process progress maps, and work plan maps that maximize user productivity.
- ▶ Workbench, which enables drag-and-drop definition of authorized configuration item space, thereby reducing time to value.

Tivoli Change and Configuration Management Database has been included into a suite of products, delivered as the IBM SmartCloud Control Desk.

For more information visit:

<http://www-01.ibm.com/software/tivoli/products/smartcloud-controldesk/>

Tivoli Application Dependency Discovery Manager

Tivoli Application Dependency Discovery Manager delivers automated discovery and configuration tracking capabilities to build application maps and provide real-time visibility into application complexity. The features of this product enable organizations to:

- ▶ Understand the structure of interdependent and complex applications
- ▶ Rapidly isolate configuration-related application problems, which reduces troubleshooting time from hours and days to minutes

- ▶ Better understand the impact of component-level events to sort issues based on application and service impact
- ▶ More effectively plan change so that application upgrades and deployments can occur without disruptions
- ▶ Create a shared topological definition of applications for use by other management applications, such as service level managers and provisioning tools

For more information:

<http://www-01.ibm.com/software/tivoli/products/taddm/>

Tivoli Provisioning Manager

Tivoli Provisioning Manager helps organizations optimize efficiency, accuracy, and service delivery by automating useful practices for data center provisioning activities, as listed here:

- ▶ Help automate best practices for common data center provisioning activities in support of change and release management processes, thereby helping to optimize efficiency, accuracy, and service delivery
- ▶ Discover and track data center resources to enable highly accurate server provisioning and software deployments
- ▶ Create thousands of virtual machines simultaneously
- ▶ Facilitate efforts to consistently follow an organization's own policies and preferred configurations, in support of corporate and regulatory compliance efforts
- ▶ Automatically provision software and configurations to Microsoft Windows servers and clients, and to Linux and UNIX servers
- ▶ Help optimize availability by maintaining configurations and managing changes to resources

For more information visit:

<http://www-01.ibm.com/software/tivoli/products/prov-mgr/>

Tivoli Service Level Advisor and Tivoli Netcool Service Quality Manager

Tivoli Service Level Advisor predicts when service level agreement violations are likely to occur, and then takes corrective actions to avoid an SLA violation.

- ▶ It enables organizations to define SLAs easily using an SLA wizard, and provide SLA evaluations as often as hourly.

- ▶ It utilizes a trend analysis algorithm that enables organizations to take a proactive approach to service level management.
- ▶ It integrates with the Tivoli Data Warehouse to consolidate and report on systems management data.
- ▶ It provides enablement of mainframe and multivendor distributed systems management data for true end-to-end service level management.

For more information visit:

<http://www-01.ibm.com/software/tivoli/products/service-level-advisor/>

Tivoli Netcool Service Quality Manager combines service level agreement and service quality management to help manage telecommunications service quality. It does this by providing end-to-end views of a service to help service providers understand quality of service from a customer's perspective, with the following key features:

- ▶ Monitoring and improving the quality of each customer experience, resulting in more effective customer care and increased customer satisfaction
- ▶ Responding to network issues based on corporate directives such as revenue, profitability, service, and customer impact
- ▶ Providing product differentiation to your enterprise sales team by offering guaranteed SLAs to attract and retain high-value enterprise customers
- ▶ Enabling the successful, rapid introduction of new services that you can offer with confidence in their service quality
- ▶ Identifying high-priority problems quickly and accurately with powerful root-cause and impact analysis
- ▶ Offering comprehensive data collection capabilities and extensive pre-established service models with full key quality indicator (KQI) and key performance indicator (KPI) mapping and reporting

For more information visit:

<http://www-01.ibm.com/software/tivoli/products/netcool-service-quality-mgr/>

Tivoli Asset Manager for IT

Tivoli Asset Manager for IT enables effective management of the IT asset lifecycle to lower cost, mitigate license compliance risk, and better align IT with business goals with the following key features:

- ▶ Optimizing software licenses and reducing the likelihood of overbuying and fines due to under-licensing

- ▶ Helping to control the cost of IT assets with a single solution that tracks and manages your hardware, software, and related information throughout their lifecycles.
- ▶ Optimizing IT asset utilization and IT service levels: deploy not more, not less
- ▶ Closely aligning IT with business requirements through IT asset cost and usage information
- ▶ Reducing time, cost, and risk associated with compliance audits through comprehensive software license management
- ▶ Relevant vendor, contract, lease, warranty and license data minimizes procurement and maintenance expenses and can help negotiate contract renewals
- ▶ Improving service desk quality and incident resolution time with accurate IT asset information

Tivoli Asset Manager for IT has been included into a suite of products, delivered as the IBM SmartCloud Control Desk.

For more information visit:

<http://www-01.ibm.com/software/tivoli/products/smartscloud-controldesk/>

VMControl and Tivoli Productivity Center

VMControl is a IBM Systems Director plug-in that decreases infrastructure costs and improves service levels by providing complete virtual server lifecycle management with the following key features:

- ▶ Brings together physical and virtual management into a single interface to reduce complexity.
- ▶ Offers unmatched cross-platform and cross-operating system management, which helps improve service delivery by eliminating isolated silos of virtualization in heterogeneous environments.
- ▶ Provides faster time-to-value and greater business agility through simplified virtualization management that allows more effective utilization of virtualized resources.
- ▶ Establishes repeatable accuracy and consistency through automation.
- ▶ Reduces operational and infrastructure costs through increased efficiency and resource utilization.

For more information visit:

<http://www-03.ibm.com/systems/software/director/vmcontrol/index.html>

Tivoli Productivity Center provides storage infrastructure management tools that can help improve time to value and reduce the complexity of managing storage environments by centralizing, simplifying and automating storage management with the following key features:

- ▶ Providing comprehensive visibility and helping centralize the management of an organization's heterogeneous storage infrastructure from a single interface using role-based administration and single sign-on
- ▶ Delivering common services for simple configuration and consistent operations across host, fabric, and storage systems
- ▶ Managing performance and connectivity from the host file system to the physical disk, including in-depth performance monitoring and analysis of the SAN fabric
- ▶ Supporting an organization's virtualization initiatives by providing insights into data usage that will help to dynamically move data across alternate tiers of storage
- ▶ Managing the capacity utilization and availability of storage systems, file systems, and databases
- ▶ Monitoring, managing, and controlling (zone) SAN fabric components
- ▶ Automating capacity provisioning of file systems
- ▶ Monitoring and tracking the performance of SAN-attached SMI-S compliant storage devices

For more information visit:

<http://www-03.ibm.com/systems/storage/software/center/>

Tivoli Service Automation Manager

Tivoli Service Automation Manager provides users the ability to request, deploy, monitor, and manage cloud computing services through traceable approval processes, with the following key features:

- ▶ Lower cost of service delivery through automation and reduced skill requirements
- ▶ Deploy IT services faster to meet the increased need for development, test, preproduction, and production systems
- ▶ Deliver a higher degree of standardization and automation for deployment and management of IT services, and reserve skilled IT staff member time for other, higher-value tasks
- ▶ Provide traceable processes and approval routings to serve as audit trails, and integrate with process governance

- ▶ Offer an integrated management capability that addresses the lifecycle changes of a cloud service
- ▶ Provide adaptable and automated preferred practices for building and managing IT infrastructures

For more information visit:

<http://www-01.ibm.com/software/tivoli/products/service-auto-mgr/>

3.3 Cloud Service Management Platform for Hybrid Cloud

This section describes the Cloud Service Management Platform components that are enhanced for Hybrid Cloud environments:

- ▶ Updated components to support Hybrid Cloud
- ▶ New components to support Hybrid Cloud

Figure 3-5 shows these components and illustrates how they map to IBM products.

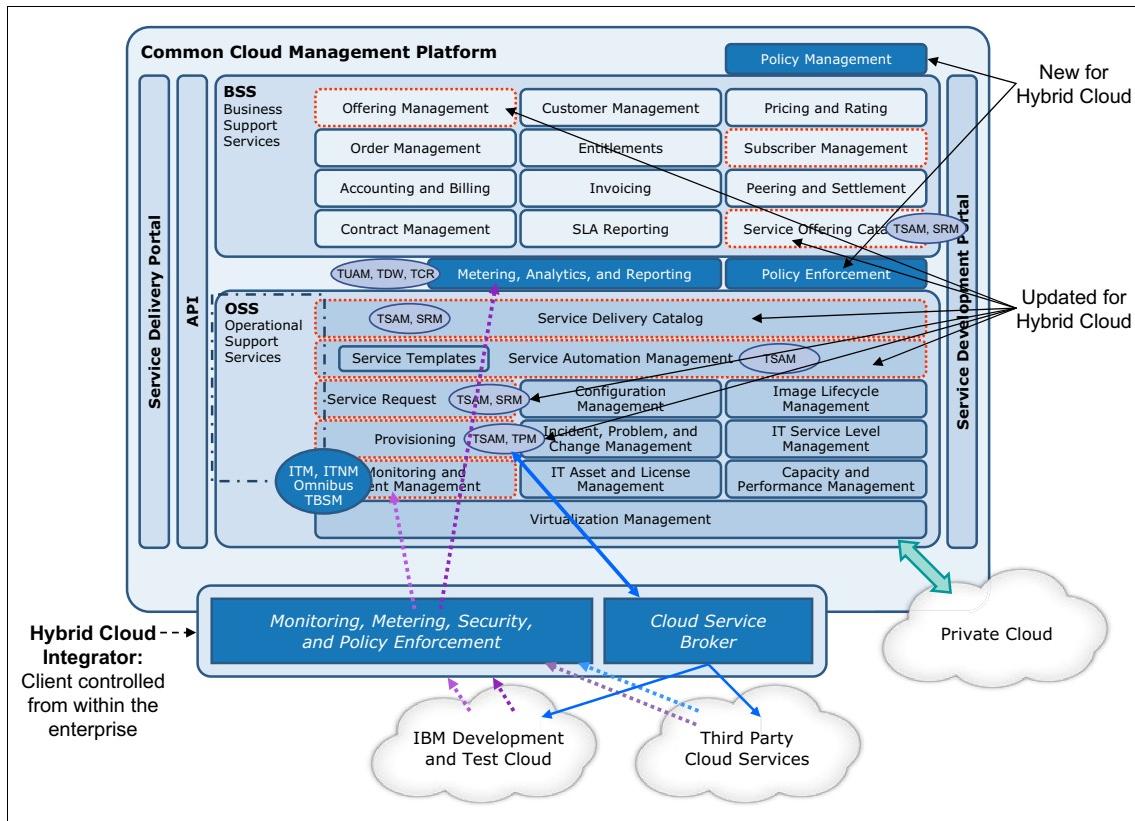


Figure 3-5 Hybrid Cloud Enhanced Service Management Platform

3.3.1 Updated components to support Hybrid Cloud

For Hybrid Cloud environments, the BSS Offering Management and Service Offering Catalog components are updated to support services that are offered both in private and public clouds.

In the OSS, the Service Delivery Catalog, Service Request, Provisioning, Service Templates and Service Automation Management components are updated to support the Hybrid Cloud environment.

These updates allow the same CCMP to extend from the private cloud space to the Hybrid Cloud space. These changes to the BSS and OSS, in conjunction with the new Policy Management and Policy Enforcement components for the cloud

service providers, allow for a seamless experience for the cloud service consumer.

3.3.2 New components to support Hybrid Cloud

Two new components, Policy Enforcement and Policy Management, are added to the BSS to support Hybrid Cloud.

Policy Management allows for a workload policy to be set for a given cloud service. For example, if a client wants a particular cloud service to never run higher than 75 percent CPU utilization, the client can set a policy for that in the Hybrid Cloud environment leveraging the BSS Policy Management features.

Next, that policy of never allowing a cloud service to use more than 75 percent CPU utilization needs a mechanism for enforcement. The Policy Enforcement component is added to the BSS to allow the enforcement of policies set in policy management. In this example, Policy Enforcement will add additional machines to a cloud service to keep CPU utilization below 75 percent so the cloud service stays within the policies set by Policy Management.



Part 2

Usage scenarios

Integrating applications across clouds can introduce unique management challenges. The Service Management Extensions for Hybrid Cloud provide the necessary features that allow companies to integrate the monitoring, provisioning, events handling, and automating of service environments located across clouds, both on and off premises.

In this part we introduce the Service Management Extensions for Hybrid Cloud and focus on the features it provides. The following chapters are included:

- ▶ Chapter 4, “Monitoring scenario” on page 61
- ▶ Chapter 5, “Provisioning scenario” on page 93
- ▶ Chapter 6, “Directory Synchronization scenario” on page 153
- ▶ Chapter 7, “Workload Governance and Management scenario” on page 173



Monitoring scenario

In Chapter 2, “Integrating cloud solutions with Cast Iron” on page 15, a scenario is introduced based on a fictional, large transnational retail company called IBM Redbooks Store with 1000 stores. This company acquires another retail company with 100 stores to add to its portfolio of products.

In this chapter we detail the monitoring portion of that scenario (2.9.2, “Monitoring” on page 35). The company plans to use its Tivoli Monitoring system to monitor CPU use to accommodate the additional capacity needed in the months of October, November, and December.

In this chapter we discuss the following topics:

- ▶ Monitoring scenario overview
- ▶ Installation and configuration
- ▶ Verifying the installation
- ▶ Running a CPU monitoring test for the Workload Governance and Management scenario
- ▶ Lessons learned
- ▶ Preferred practices
- ▶ Summary

4.1 Monitoring scenario overview

The main objective of the monitoring scenario is to provide the ability to monitor servers deployed in the public clouds using an on-premise deployment of IBM Tivoli Monitoring. Virtual machines in the cloud are monitored and results are displayed along with the on-premise machines.

Monitoring in a hybrid cloud environment serves the same purposes as it does in the traditional enterprise: administrators can monitor availability, resource usage, problems, and service-level compliance. In addition to visibility, monitoring provides automation and notification capabilities.

Monitoring has a special importance in a hybrid environment. When the computing environment capacity can dynamically increase or decrease, administrators need visibility. Increased capacity and performance are balanced against increased cost. At one extreme, a solution is too costly if utilization is low. At the other extreme, performance can suffer if the resources are insufficient. Administrators need visibility and usage data to make decisions about the right-sized environment for their needs.

The Monitoring scenario utilizes extensions to IBM Tivoli Monitoring (ITM), providing the ability to monitor servers within IBM SmartCloud and Amazon EC2. In the example shown in Figure 4-1 on page 63, IBM SmartCloud is used.

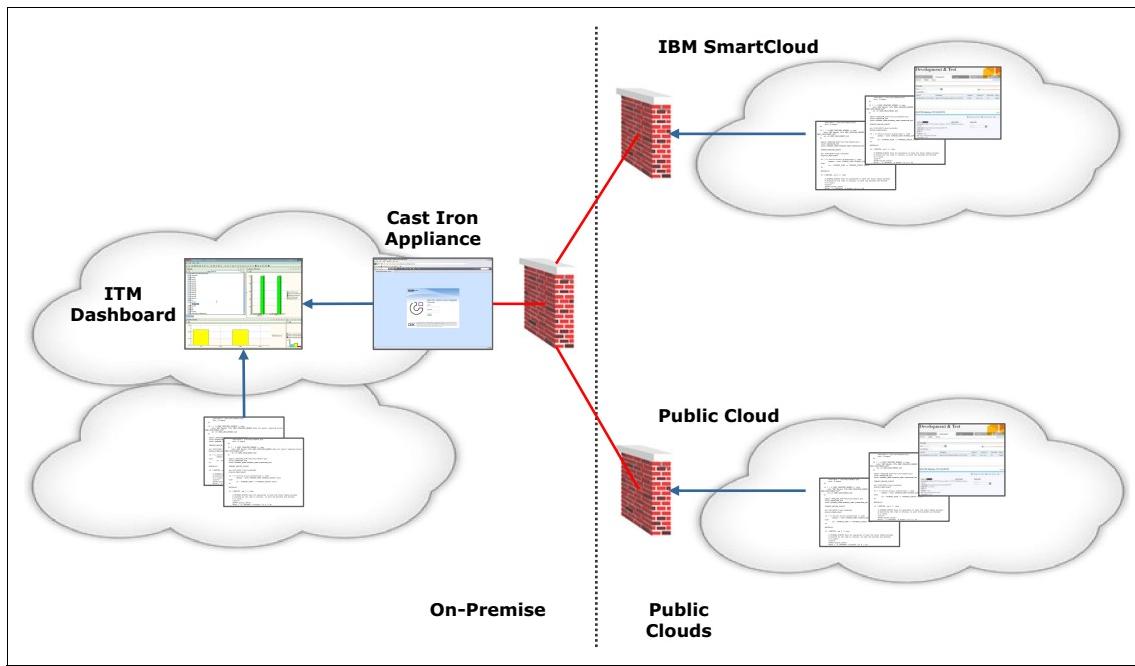


Figure 4-1 Monitoring scenario at a glance

IBM Tivoli Monitoring resides on premise. A cloud gateway agent is installed and configured. It allows traffic to pass through a firewall. Agents are deployed onto the virtual machines in the cloud.

In this scenario, the IBM Redbooks store experiences peak seasonal usage each year in October through December. The system is designed to provision into the cloud for peak needs. The requirement is that both the machines in the cloud and the on-premise machines need to be monitored. In particular, the CPU usage of the cloud machines are monitored to detect high utilization rates. High utilization rates trigger alerts that can be automated to provision additional resources. The key role monitoring plays is to detect when the elastic cloud needs expand.

4.2 Installation and configuration

This section describes how to set up the monitoring scenario. It does not address the full installation of the products. Only the installation steps that are relevant for the integration monitoring scenario are discussed.

4.2.1 Install Service Management Extensions for Hybrid Cloud plug-ins

This section explains the high level steps needed to install the Service Management Extensions for Hybrid Cloud plug-ins. Plug-ins add the integration capabilities needed to pass control and data between the monitoring server that is on premise and the agents in the cloud.

1. The WebSphere Cast Iron integration appliance must be at revision level 6.1.0.2 or higher. To verify the revision level, log on the Web Management Console using a supported browser. Use the management IP address for the integration appliance.

<http://CIname/login/>

- From the Web Management Console, select **System** and verify the operating system level as shown in Figure 4-2.

The screenshot shows a Mozilla Firefox browser window titled "WebSphere Cast Iron web management console - Mozilla Firefox: IBM Edition". The address bar displays "http://10.7.28.200/". The main content area is titled "WebSphere Cast Iron web management console" and shows the "System Summary" page. On the left, there is a navigation sidebar with links: Home, Repository, Logs, Security, Network, System (which is selected and highlighted in blue), Upgrade, Staging DB, Hardware Status, and Commands. The "System Summary" table contains the following information:

| System Summary | |
|----------------|--|
| Model | Cast Iron vA3000, Revision A |
| Version | Cast Iron Operating System 6.1.0.2.0000 (Mon Oct 17 12:38:38 UTC 2011) |
| Serial Number | VMWYF7VDBMKP4X5K |

Below the table are two buttons: "Upgrade Operating System" and "Update Connector Libraries".

Figure 4-2 System summary in the Web Management Console

- Log on to the WebSphere Cast Iron Design Console at:

<http://CIname/express/>

CIname is the host name or management IP address of the WebSphere Cast Iron integration appliance. The default values are:

- Username: admin
- Password: !n0r1t5@C

4. Click the user name in the top right corner of the WebSphere Cast Iron Design Console application window (Administrator in Figure 4-3).

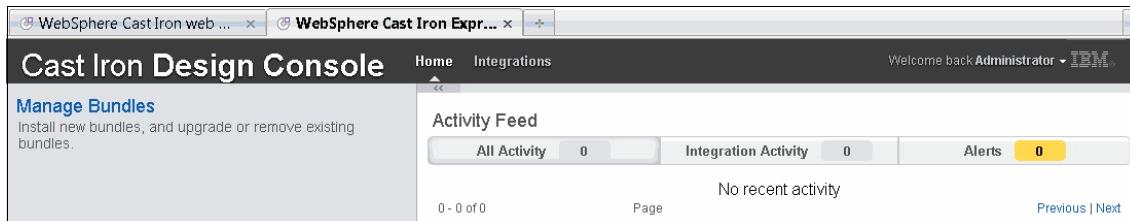


Figure 4-3 WebSphere Cast Iron Design Console

5. From the drop-down menu, select **Manage Bundles** from the list or in the Cast Iron Design Console window, then click **Manage Bundles** on the left navigation panel to open the panel shown in Figure 4-4. Then click **Install New Bundle**.

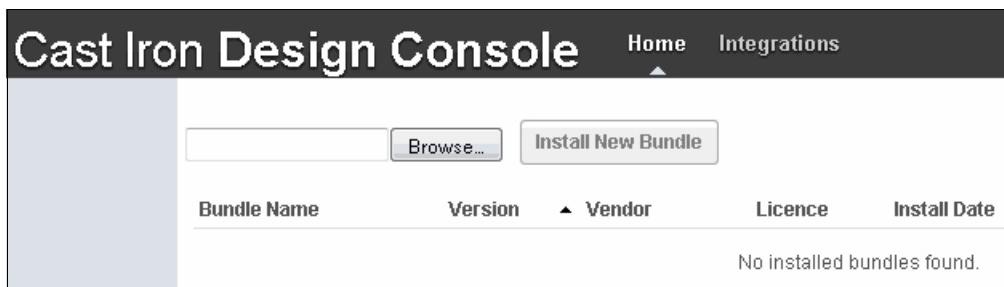


Figure 4-4 Manage bundles

6. Click **Browse** and select a .brcypt bundle file (Figure 4-5 on page 67). Notice there are different bundles for a physical or virtual WebSphere Cast Iron Appliance.

Install the *monitoring-bundle_virtual.bcrypt* plug-in file if you are using the virtual appliance, WebSphere Cast Iron Hypervisor Edition.

Install the *monitoring-bundle_physical.bcrypt* plug-in file if you are using the physical hardware appliance, WebSphere DataPower Cast Iron Appliance XH40.

Then click the **Upload** button. In this scenario, the hypervisor edition is used.

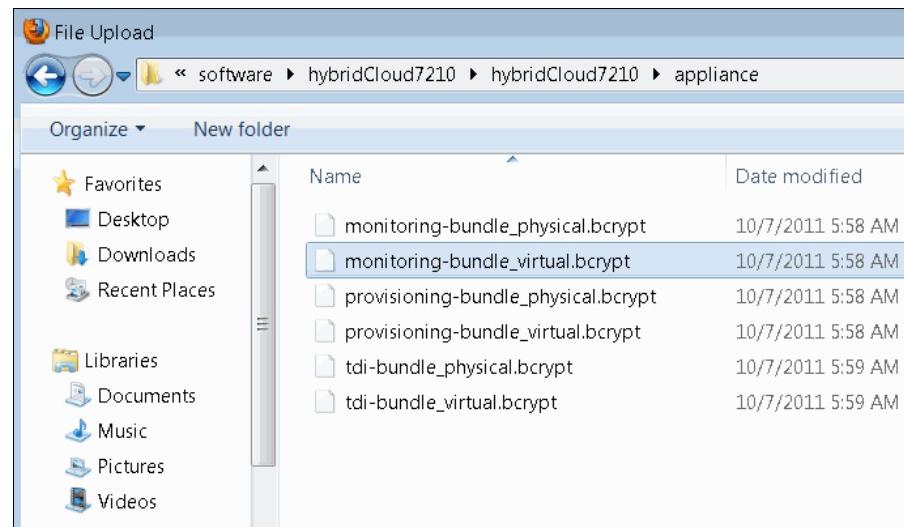


Figure 4-5 Manage bundles installation list

7. Click **Install New Bundle** (Figure 4-6).

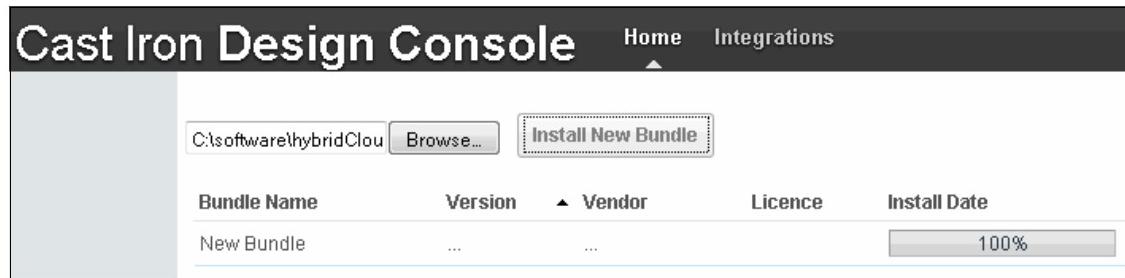


Figure 4-6 Install the new bundle

8. The License Agreement is displayed. To install the plug-in, select **Accept**.

9. Verify that the installation was successful (Figure 4-7) and click **OK**.

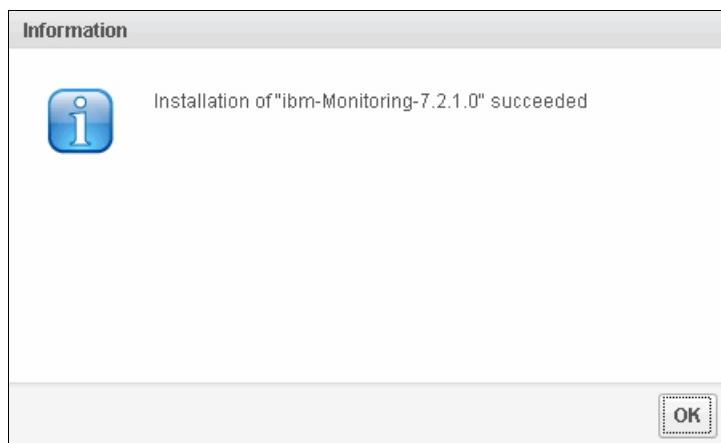


Figure 4-7 Checking the installation status

10. Figure 4-8 shows that the Service Management Extensions for Hybrid Cloud monitoring plug-in was successfully installed.

A screenshot of the "Cast Iron Design Console" interface. The top navigation bar includes "Home" and "Integrations". Below the header, there are two buttons: "Browse..." and "Install New Bundle". A table displays the details of a successfully installed bundle:

| Bundle Name | Version | Vendor | Licence | Install Date |
|-------------|---------|--------|---------|---------------------|
| Monitoring | 7.2.1.0 | ibm | File | 2012-03-01 21:54:45 |

Figure 4-8 Version, date, and time of successful installation monitoring plug-in

4.2.2 Install IBM Cloud Gateway Agent on IBM SmartCloud Enterprise

This section explains the function of the IBM Tivoli Monitoring Cloud Gateway Agent and its installation.

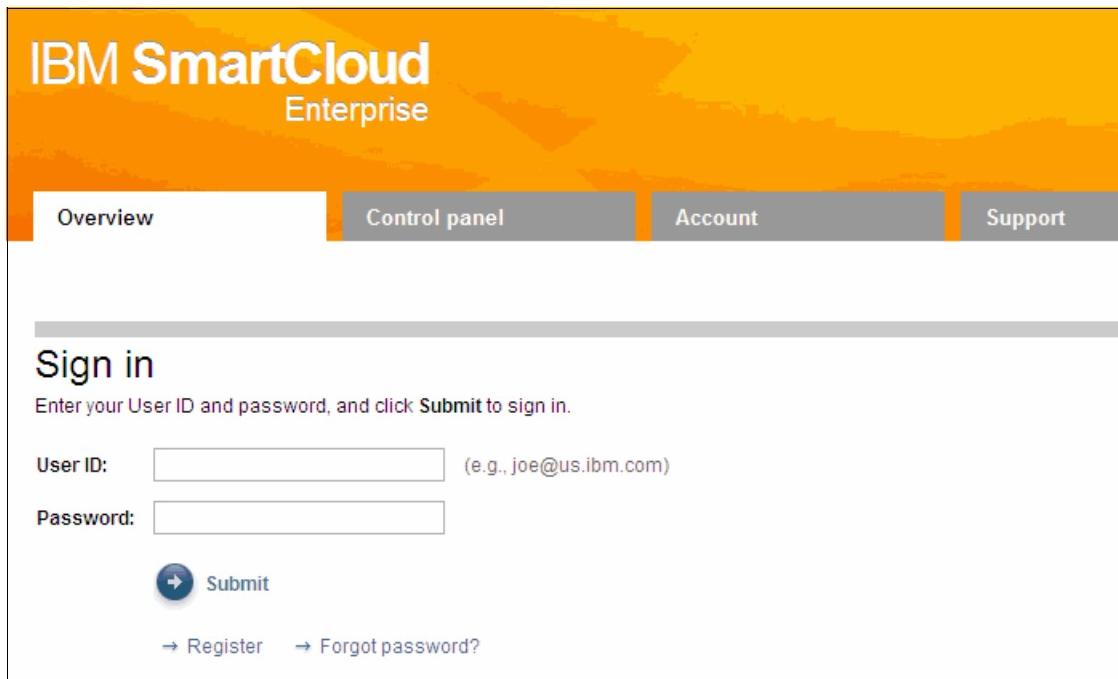
The IBM Tivoli Monitoring Cloud Gateway Agent acts as a proxy for the Tivoli Enterprise Monitoring Server and listens for connections from the enterprise into the cloud. The Cloud Gateway Agent integrates the resource metrics for other cloud agents. This agent is required for the monitoring and workload governance scenarios.

In this scenario, the IBM SmartCloud is used. An account on the cloud system is a prerequisite.

To install the Cloud Gateway Agent, follow these steps:

1. Create an instance of an image in the IBM SmartCloud Enterprise.

For more information, see the IBM SmartCloud Enterprise User's Guide at
<https://www-147.ibm.com/cloud/enterprise/dashboard>



The image shows the IBM SmartCloud Enterprise welcome panel. At the top, there is a navigation bar with four tabs: Overview, Control panel (which is highlighted in blue), Account, and Support. Below the navigation bar, the page title is "IBM SmartCloud Enterprise". The main content area is titled "Sign in" and contains the following text: "Enter your User ID and password, and click Submit to sign in." There are two input fields: "User ID:" followed by a text input box containing "(e.g., joe@us.ibm.com)" and "Password:" followed by a text input box. Below the input fields is a "Submit" button with a circular arrow icon. At the bottom of the form, there are links for "Register" and "Forgot password?".

Figure 4-9 IBM SmartCloud Enterprise welcome panel

2. Install the Cloud Gateway Agent.

In this step, a Linux Tivoli Monitoring agent in the IBM SmartCloud Enterprise cloud is installed for use as a gateway:

- a. Log on the Linux instance at the IBM SmartCloud Enterprise environment using the default account *idcuser*.
- b. Start a root shell session with the **sudo bash** command.

- c. Install the Linux Tivoli Monitoring agent into the instance, as described in
http://publib.boulder.ibm.com/infocenter/tivihelp/v15r1/topic/com.ibm.itm.doc_6.2.2fp2/installation.htm#installation

Example 4-1 shows a typical installation sequence using the `./install.sh` command.

Example 4-1 Installation of the Linux Tivoli Monitoring agent

```
./install.sh
```

```
Enter the name of the IBM Tivoli Monitoring directory  
[ default = /opt/IBM/ITM ]:
```

```
ITM home directory "/opt/IBM/ITM" already exists.  
OK to use it [ 1-yes, 2-no; "1" is default ]?
```

```
Select one of the following:
```

- 1) Install products to the local host.
- 2) Install products to depot for remote deployment (requires TEMS).
- 3) Install TEMS support for remote seeding
- 4) Exit install.

```
Please enter a valid number: 1
```

```
The following products are available for installation:
```

- 1) Agentless Monitoring for AIX Operating Systems V06.22.02.00
- 2) Agentless Monitoring for HP-UX Operating Systems V06.22.02.00
- 3) Agentless Monitoring for Linux Operating Systems V06.22.02.00
- 4) Agentless Monitoring for Solaris Operating Systems V06.22.02.00
- 5) Agentless Monitoring for Windows Operating Systems V06.22.02.00
- 6) Monitoring Agent for Linux OS V06.22.02.00
- 7) Monitoring Agent for UNIX Logs V06.22.02.00
- 8) Tivoli Enterprise Services User Interface Extensions V06.22.02.00
- 9) Universal Agent V06.22.02.00
- 10) all of the above

Type the numbers for the products you want to install, type "b" to change operating system, or type "q" to quit selection.
If you enter more than one number, separate the numbers by a comma or a space.

Type your selections here: 6

The following products will be installed:

Monitoring Agent for Linux OS V06.22.02.00

Are your selections correct [1=Yes, 2=No ; default is "1"] ?

... installing "Monitoring Agent for Linux OS V06.22.02.00 for Linux Intel R2.6 (32 bit)"; please wait.

=> installed "Monitoring Agent for Linux OS V06.22.02.00 for Linux Intel R2.6 (32 bit)".

... Initializing component Monitoring Agent for Linux OS V06.22.02.00 for Linux Intel R2.6 (32 bit).

... Monitoring Agent for Linux OS V06.22.02.00 for Linux Intel R2.6 (32 bit) initialized.

Do you want to install additional products or product support packages [1=Yes, 2=No ; default is "2"] ?

... postprocessing; please wait.

... finished postprocessing.

Installation step complete.

You may now configure any locally installed IBM Tivoli Monitoring product via the "/opt/IBM/ITM/bin/itmcmd config" command.

3. Configure the Cloud Gateway Agent.

This step will configure the Cloud Gateway Agent in the IBM SmartCloud Enterprise cloud to connect to the local IBM Tivoli Enterprise Monitoring Server:

- a. Log on to the Linux IBM SmartCloud Enterprise instance as the default user idcuser, as described in *IBM SmartCloud Enterprise User's Guide*.
- b. Use the **sudo bash** command to start a shell session as the root user.
- c. Enter the following command:

```
# /opt/IBM/ITM/bin/itmcmd config -A 1z
```

- d. Accept the default values at each prompt, *except* for the TEMS Host Name prompt.

When prompted for the TEMS Host Name, enter `localhost` to ensure that the configuration remains independent of the actual host name. The agent connects to itself because it acts as the cloud gateway, and provides monitoring information about the host where it is installed.

```
Will this agent connect to a TEMS? [1=YES, 2=No] (Default is: 1):
TEMS Host Name (Default is: vhost1165): localhost
```

- e. Create the `/opt/IBM/ITM/gateway` directory and change to the gateway directory using the following commands:

```
# mkdir /opt/IBM/ITM/gateway
# cd /opt/IBM/ITM/gateway
```

- f. Configure the Cloud Gateway Agent firewall gateway:

- i. Create the `/opt/IBM/ITM/gateway/gateway.xml` file.

- ii. Copy and paste the information shown in Example 4-2 to the file.

Example 4-2 Gateway.xml contents

```
<tep:gateway xmlns:tep="http://xml.schemas.ibm.com/tivoli/tep/kde/" name="sproxy"
threads="32">
  <zone name="DMZ2" maxconn="512" error="ignore">
    <interface name="uprelay" role="listen">
      <!-- Let the on-premises gateway server talk to us over port 10,005. -->
      <bind localport="10005"/>
      <!--
          Ports that we proxy over to the on-premises gateway.
          Important: The "service" tag values here must be an exact match with
          the services defined in the on-premises side. Keep them in synch.
      -->
      <interface name="serverproxy" role="proxy">
        <!-- IP.PIPE Port 1918 is for TEMS, 6014 is for Warehouse Proxy (TDW) -->
        <bind localport="1918" service="tems_pipe" />
        <bind localport="6014" service="whp_pipe" />
        <bind localport="1918" service="tems_pipe" ipversion="6"/>
        <bind localport="6014" service="whp_pipe" ipversion="6"/>
        <!--IP.SPIPE Port 3660 is for TEMS, 7756 is for Warehouse Proxy (TDW)-->
        <bind localport="3660" service="tems_spipe" />
        <bind localport="7756" service="whp_spipe" />
        <bind localport="3660" service="tems_spipe" ipversion="6"/>
        <bind localport="7756" service="whp_spipe" ipversion="6"/>
      </interface>
    </interface>
  </zone>
```

```
</tep:gateway>
```

- g. Change to the /opt/IBM/ITM/config directory.

```
# cd /opt/IBM/ITM/config
```

- h. Edit the 1z.ini configuration file. Add or modify the following three attributes and values:

```
CTIRA_HEARTBEAT=5  
CTIRA_RECONNECT_WAIT=5  
KDE_GATEWAY=/opt/IBM/ITM/gateway/gateway.xml
```

CTIRA_HEARTBEAT and CTIRA_RECONNECT_WAIT values are in minutes. In this case CTIRA_HEARTBEAT=5 indicates five minutes between heartbeats.

Warning: Modifying these attributes in this way for every agent can impact performance.

- i. Ensure the firewall ports specified in the gateway.xml file are open.

The gateway.xml file contains a localport attribute. Use this attribute to specify inbound ports to allow the Cloud Gateway Agent to connect to the Tivoli Enterprise Monitoring Server. You also need to specify a port of your choice to connect the Enterprise Gateway Agent to the Cloud Gateway Agent.

These ports must be open for inbound connections on operating systems where the agent is installed.

If the Tivoli Enterprise Monitoring Server uses a IP.PIPE connection, open port 1918 (Tivoli Enterprise Monitoring Server) and port 6014 (Warehouse Proxy agent). For an IP.SPIPE connection, open port 3660 (Tivoli Enterprise Monitoring Server) and port 7756 (Warehouse Proxy agent). You will also need port 10005 opened for inbound connections to enable the enterprise gateway agent to connect to the cloud gateway agent.

4. Save the image for reuse. You can use this instance as a template to provision multiple cloud images.
5. Start the Cloud Gateway Agent.

To start the IBM Cloud Gateway agent use **itmcmd** command, as follows:

```
# cd /opt/IBM/ITM/bin  
# ./itmcmd agent start 1z  
Starting Monitoring Agent for Linux OS ...  
Monitoring Agent for Linux OS started
```

6. Configure Cloud Agents to connect to the Cloud Gateway Agent.
 - a. Use the same process to install the ITM Linux Monitoring Agent.
 - b. After installation is complete, using the **sudo bash** command to start a shell session as the root user, and run the following command:
`# /opt/IBM/ITM/bin/itmcmd config -A 1z`
 - c. Accept all the default value at each prompt *except* for the TEMS Host Name prompt.

When prompted for the TEMS Host Name, enter the IP address of the Cloud Gateway machine. The agent connects to the cloud gateway, which forwards monitoring information to the TEMS in the private cloud.

Will this agent connect to a TEMS? [1=YES, 2=No] (Default is: 1):
TEMS Host Name (Default is: vhost1165): **170.225.100.145**

7. During the Linux IBM Tivoli Monitoring agent install, the `/etc/init.d/ITMAgent1` file is generated. This process enables an autostart of the Tivoli Monitoring agents as the root user. However, by default, the root user does not have a login shell in the cloud.

Modifications to the file will be overwritten, and will need to be repeated, if the `/opt/IBM/ITM/bin/itmcmd config -A 1z` configuration command is executed again.

To enable the root user to autostart cloud agents, make the following modifications to the `ITMAgent1` file.

The original entry is shown in Example 4-3.

Example 4-3 Original ITMAgent1 commands

```
start_all()
{
/bin/su - root -c " /opt/IBM/ITM/bin/itmcmd agent start 1z
>/dev/null 2>&1"
}
stop_all()
{
/bin/su - root -c " /opt/IBM/ITM/bin/itmcmd agent stop 1z >/dev/null
2>&1"
}
```

The edited entry is shown in Example 4-4.

Example 4-4 Edited ITMAgent1 commands

```
start_all()
{
```

```
sudo $CANDLEHOME/bin/itmcmd agent start lz >/var/log/itmcmd.log 2>&1
}
stop_all()
{
sudo $CANDLEHOME/bin/itmcmd agent stop lz >/dev/null 2>&1
}
```

4.2.3 Creating the integration

In this section we explain how to create an integration in the Cast Iron appliance to monitor Linux on public clouds:

1. Log on to the WebSphere Cast Iron Design Console.
2. Select the source and target of the integration:
 - a. Click **Integrations** on the main window.
 - b. Select **Tivoli Monitoring Server** as the source.
 - c. Select **Off Premise Agent** as the target.

Figure 4-10 shows the results.

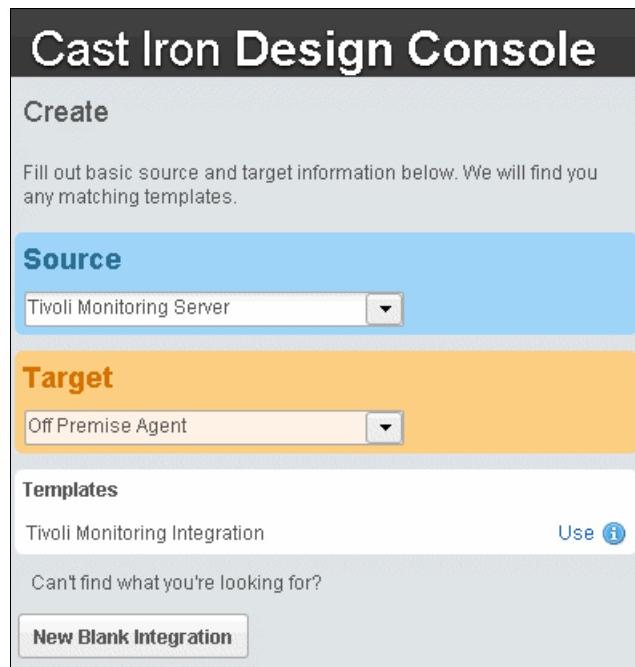


Figure 4-10 Select the source and target of the integration

- d. Click **New Blank Integration**.

Tip: After installing the monitoring Service Management Extension for Hybrid Cloud plug-in, if it does not display Tivoli Monitoring Server at the drop-down box, then refresh your browser.

3. Enter a unique and descriptive name for the integration and save. In this example, the name is Tivoli Monitoring Server to Off Premise Agent.

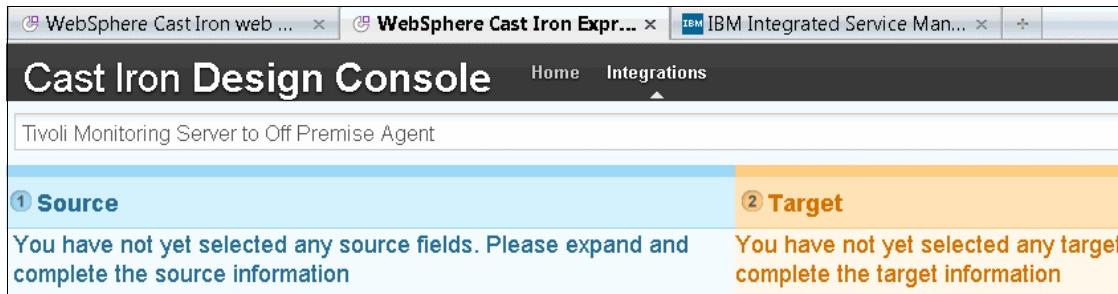


Figure 4-11 Unique name for integration

4. Click the Source tab to enter the connection properties for the source; see Figure 4-12 on page 77.
 - a. Specify a Name and Hostname (or IP Address)
 - b. Verify the port configuration and protocol IP.

Cast Iron Design Console

Home Integrations ▾

Tivoli Monitoring Server to Off Premise Agent

① Source

Connection Type

Tivoli Monitoring Server

Existing Connections

▼ Details

Name *

TEMS Source

Hostname *

192.168.1.104

Port *

1918

Warehouse port *

6014

IP.PIPE/IP.SPIPE *

ip.pipe

Connect

Figure 4-12 Source configuration

The default values are:

- Port 1918 for TEMS
 - Port 6014 for ITM Warehouse
 - Protocol IP is ip.pipe
- c. Click **Connect** to start the connection; see Figure 4-13 on page 78.

The screenshot shows the Cast Iron Design Console interface. At the top, there's a navigation bar with 'Home' and 'Integrations'. Below it, a title bar says 'Tivoli Monitoring Server to Off Premise Agent'. The main area is divided into two columns. The left column, titled '① Source', contains a yellow message box saying 'CTGHC1029I: Connection information accepted.' It also includes fields for 'Connection Type' (set to 'Tivoli Monitoring Server'), 'Existing Connections' (empty), and a 'Details' section with fields for 'Name' ('TEMPS Source'), 'Hostname' ('192.168.1.104'), 'Port' ('1918'), and 'Warehouse port' ('6014'). The right column has sections for 'Operations' (set to 'None selected') and 'Available Objects' and 'Available Fields' tables, which are currently empty.

Figure 4-13 Source connection accepted

5. Click the **Target** tab to enter the connection properties for the target of the integration (Figure 4-14 on page 79):
 - a. Specify a Name and Hostname (or IP Address).
 - b. Specify the port number for the outbound connection from the appliance to the off-premise agent.
 - c. Click **Connect**; see Figure 4-14 on page 79.

Cast Iron Design Console

Home Integrations

Tivoli Monitoring Server to Off Premise Agent

② Target

CTGHC1029I: Connection information accepted.

Connection Type: Off Premise Agent

Existing Connections:

▼ Details

Name: Off_Premiss_lz_agent

Hostname: 170.225.100.145

Port: 10005

Operations: None selected

Available Objects

Connect

The screenshot shows the 'Target' configuration screen in the Cast Iron Design Console. A green checkmark icon indicates that the connection information has been accepted. The 'Connection Type' is set to 'Off Premise Agent'. In the 'Details' section, the 'Name' is 'Off_Premiss_lz_agent', 'Hostname' is '170.225.100.145', and 'Port' is '10005'. The 'Operations' dropdown is set to 'None selected'. On the right side, there is a 'Available Objects' panel with a tree view. At the bottom, there is a 'Connect' button.

Figure 4-14 Target connection accepted

Use the same port: This port must be the same as configured in gateway.xml on the agent.

```
<!-- Let the on-premises gateway server talk to us over port 10,005. -->
<bind localport="10005"/>
<!--
```

Our example is 10005.

6. Return to the original window and click **Save** to save the integration (Figure 4-15).



Figure 4-15 Save integration

A message is displayed indicating the integration has been saved successfully (Figure 4-16).

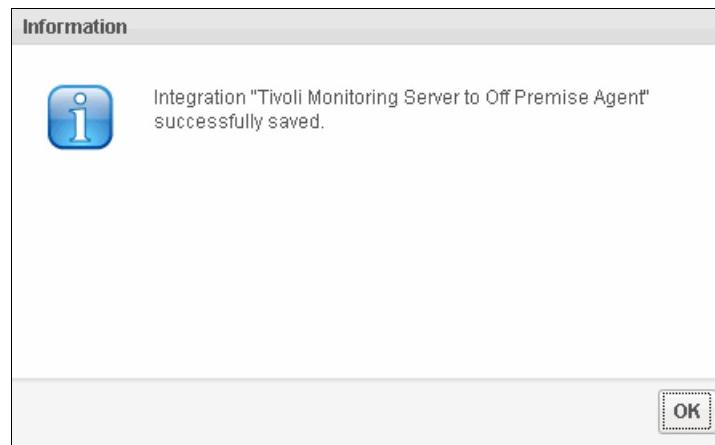


Figure 4-16 Successful integration saved

7. To start the integration, click the play icon (Figure 4-17).

| Name | Updated | Updated by | Status |
|---|-------------------------------|------------|---------|
| TSAM to IBM Compute Cloud | March 2, 2012 12:44:34 AM GMT | admin | Running |
| Tivoli Monitoring Server to Off Premise Agent | March 9, 2012 10:29:57 PM GMT | admin | Stopped |

Figure 4-17 Start integration

8. Verify that the integration is running successfully (Figure 4-18).

| Name | Updated | Updated by | Status |
|--|-------------------------------|------------|---------|
| ITSO Tivoli Monitoring Server to Off Premise Agent | March 14, 2012 1:46:22 PM GMT | andre | Running |

Figure 4-18 Integration started successfully

4.3 Verifying the installation

Perform the following steps to check monitoring status at server on the IBM SmartCloud Enterprise.

4.3.1 Log in to IBM Tivoli Monitoring

There are several ways to connect to IBM Tivoli Monitoring. In this example we use the web application.

1. Open the Tivoli Enterprise Portal Web page:

`http://TEP address:1920///cnp/kdh/lib/cnp.html`

2. Use your login and password to connect, as shown in Figure 4-19.

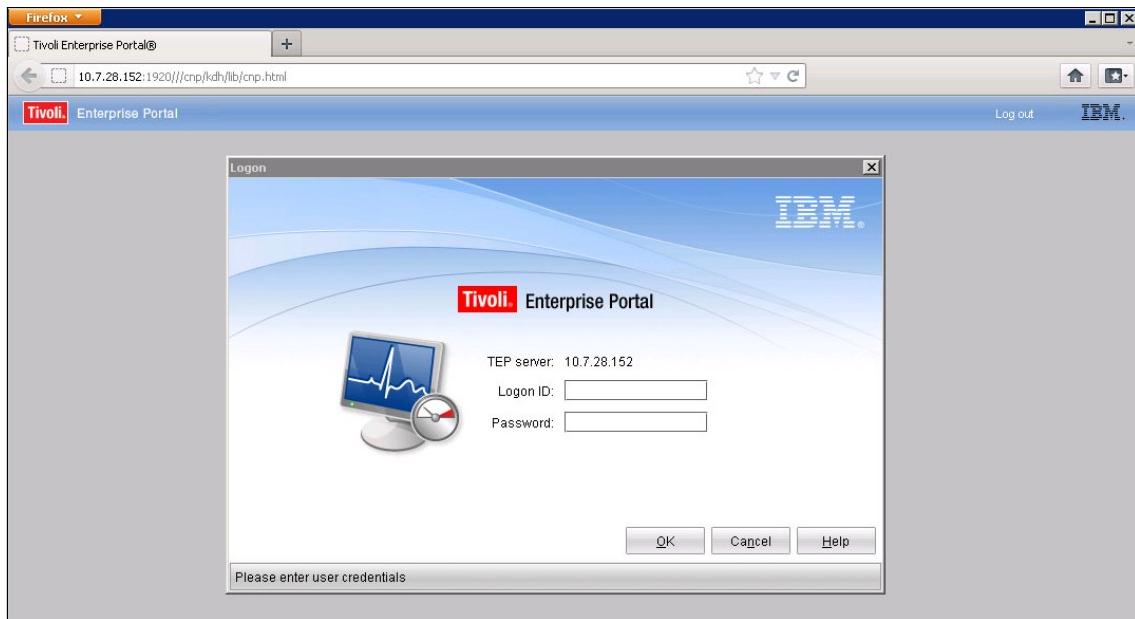


Figure 4-19 TEP Login panel

4.3.2 Verifying the new Linux agent is shown in dashboard

Figure 4-20 shows the Tivoli Enterprise Portal dashboard before the monitoring integration has been started.

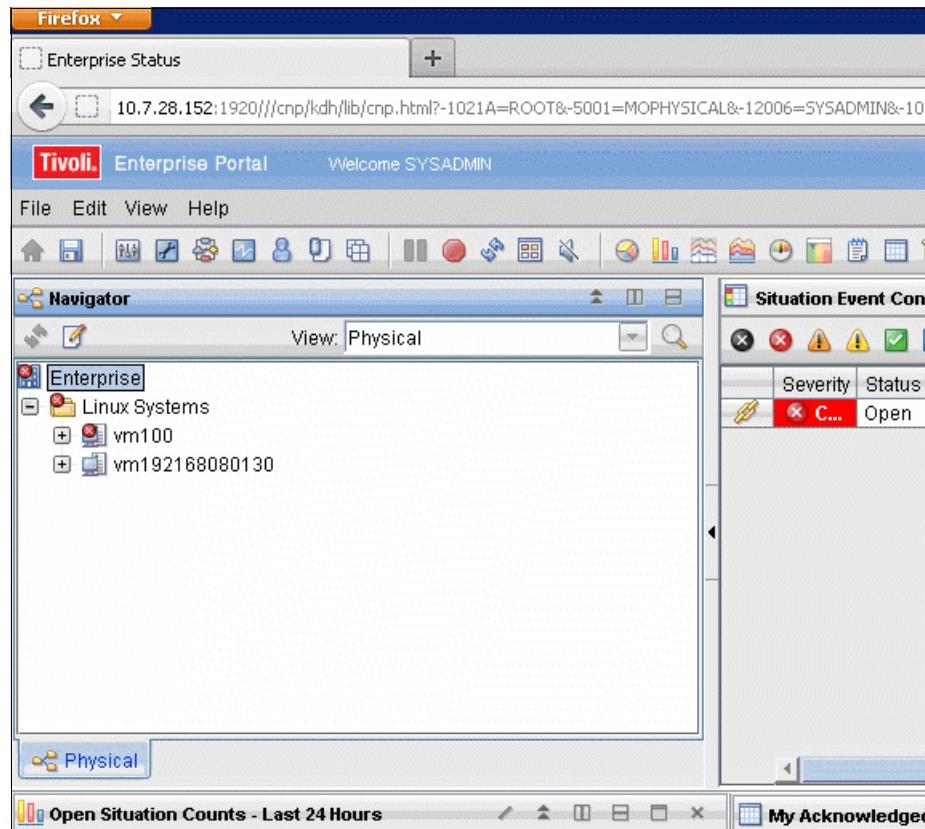


Figure 4-20 Initial TEP dashboard

After the integration is up and running, the SmartCloud Enterprise Linux agent appears in the Tivoli Enterprise Portal; see Figure 4-21.

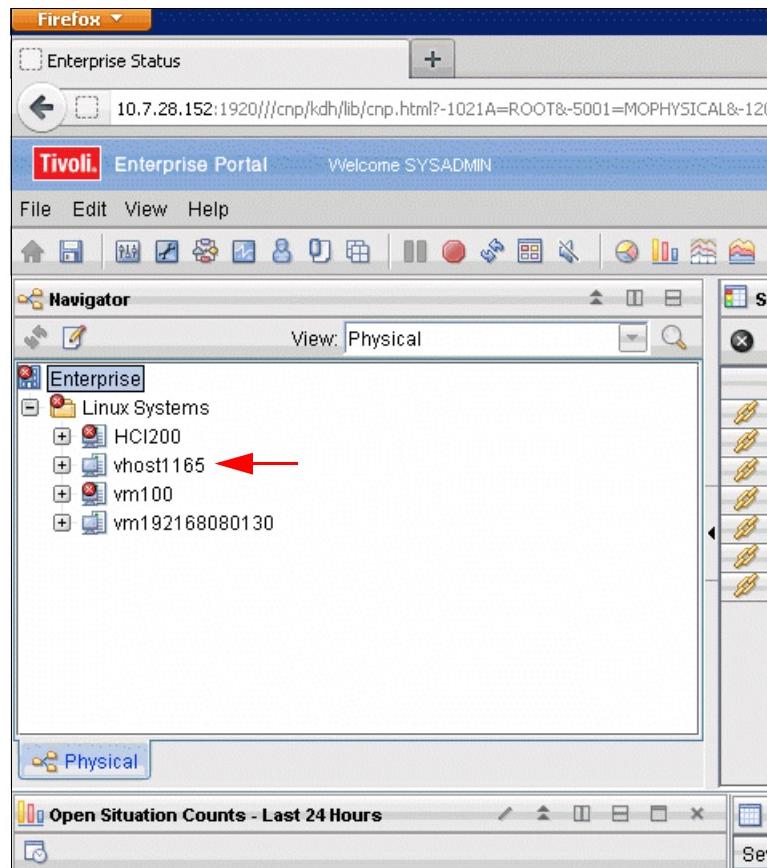


Figure 4-21 TEP Dashboard after integration

4.3.3 Verifying the monitoring scenario is working

Perform the following steps to create a situation monitoring test:

1. On the IBM Tivoli Monitoring dashboard, create a *situation*. A situation raises an alert when the monitored condition occurs. Situations can be used to alert operations staff, but they can also be forwarded to an event handler to trigger automated responses.

Using the right-click button, click the **Linux server** icon and select **situations** to open the situation editor.

2. Click the second icon at the top of the window to create a new situation.

- To verify that monitoring is working properly, create a ping test situation. This situation, shown in Figure 4-22, raises an alert if the ping process is not running on the monitored agent.

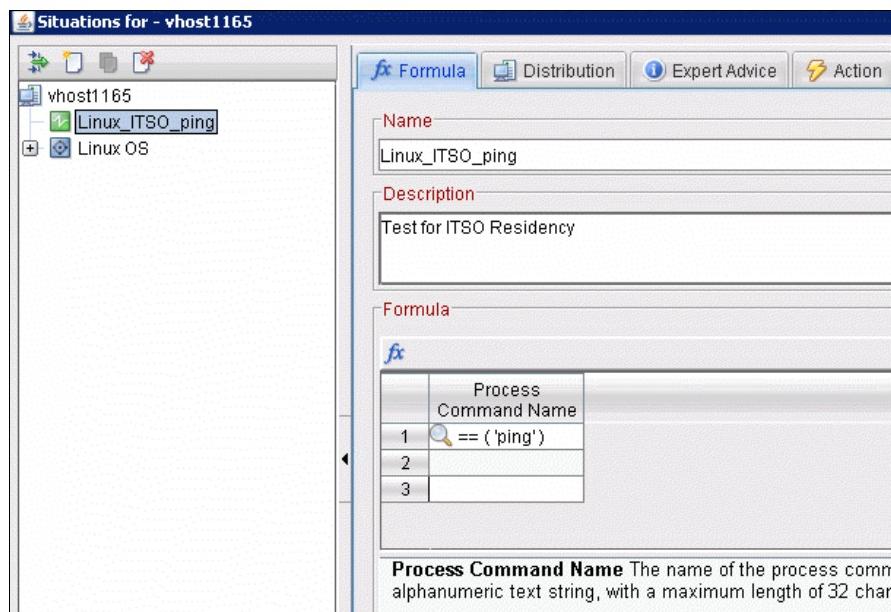


Figure 4-22 Ping test situation

4. Apply the situation and check the ITM dashboard. On the monitored system, verify that the ping process does not exist. Then verify that an alert is triggered as shown in Figure 4-23.

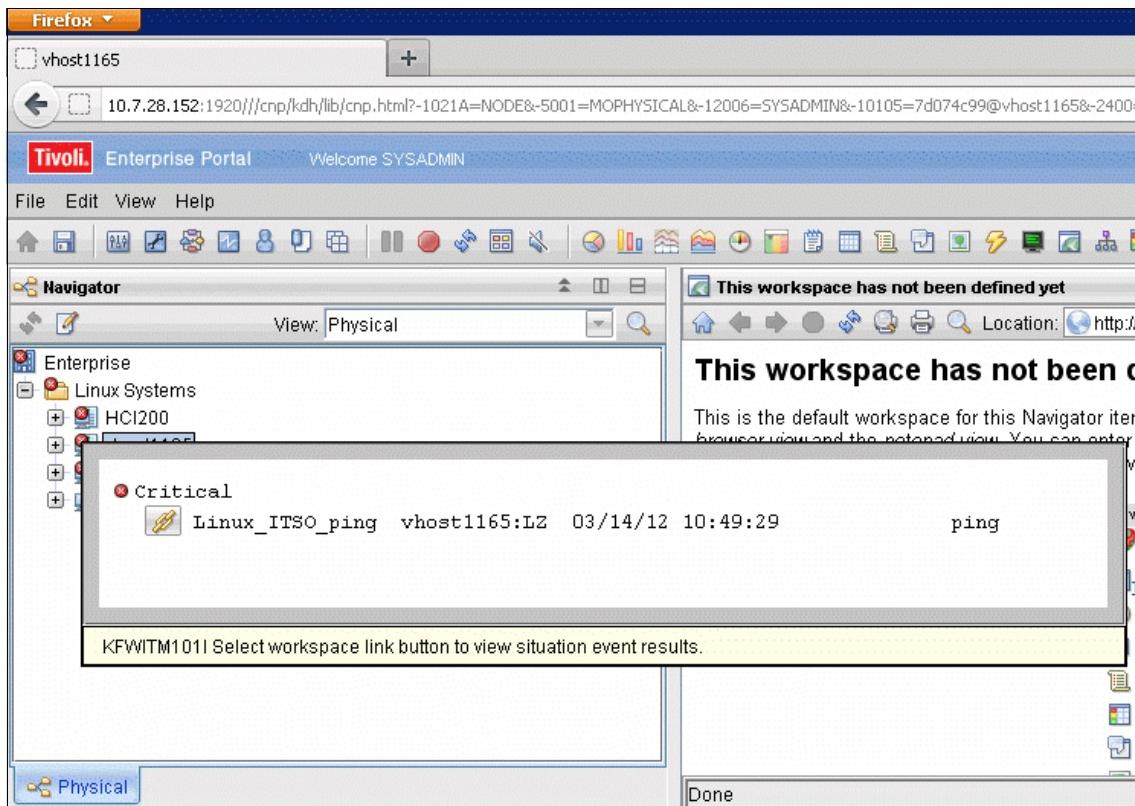


Figure 4-23 Dashboard after enable ping situation

5. This test verifies that the integration between IBM Tivoli Monitoring and WebSphere Cast Iron is up and running. Just to double-check, create the ping process by running ping on the monitored system as shown in Figure 4-24.

```
[root@vhost1165 bin]# ping www.redbooks.ibm.com
PING dispasd-45-redbk.boulder.ibm.com (170.225.15.45) 56(84) bytes of data.
64 bytes from 170.225.15.45: icmp_seq=1 ttl=242 time=7.75 ms
64 bytes from 170.225.15.45: icmp_seq=2 ttl=242 time=7.53 ms
64 bytes from 170.225.15.45: icmp_seq=3 ttl=242 time=7.85 ms
64 bytes from 170.225.15.45: icmp_seq=4 ttl=242 time=7.85 ms
```

Figure 4-24 Ping process running

6. Close the existing alert so the dashboard now looks like Figure 4-25.

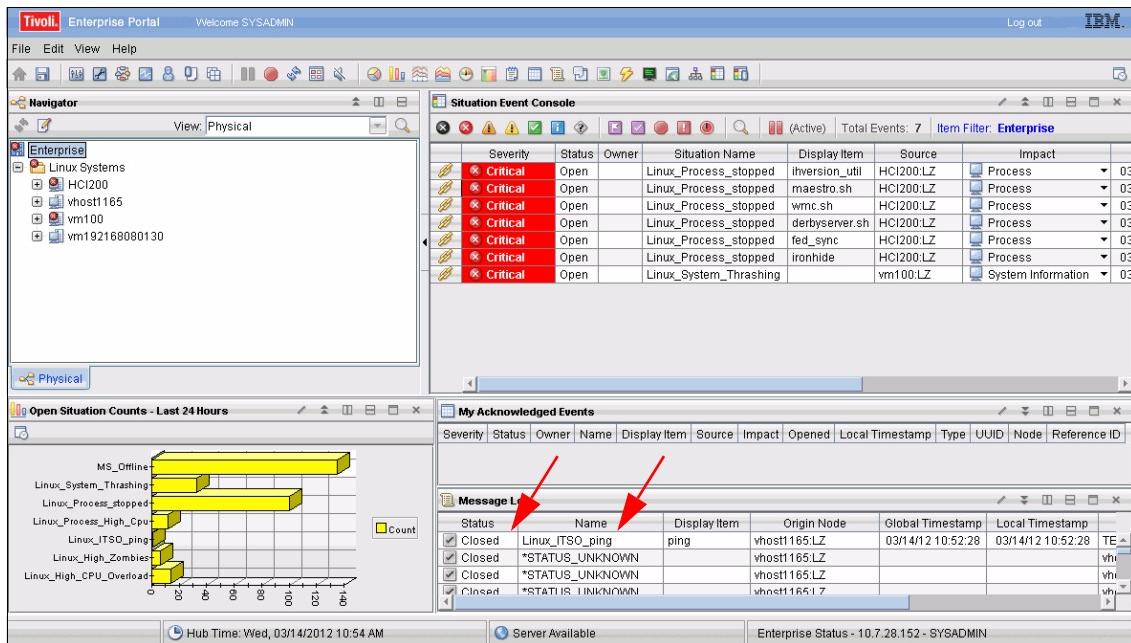


Figure 4-25 Dashboard after ping process start

7. Check the message log to verify the alert status is closed (Figure 4-26).

| Severity | Status | Owner | Name | Display Item | Source | Impact | Opened | Local Timestamp | Type | UUID | Node | Reference ID |
|-------------|--|-------|-----------------|--------------|--------------|--------|-------------------|-------------------|------|------|------|--------------|
| Message Log | | | | | | | | | | | | |
| | <input checked="" type="checkbox"/> Closed | | Linux_ITSO_ping | ping | vhost1165:LZ | | 03/14/12 10:52:28 | 03/14/12 10:52:28 | TE | | | |
| | <input checked="" type="checkbox"/> Closed | | *STATUS_UNKNOWN | | vhost1165:LZ | | | | vhi | | | |
| | <input checked="" type="checkbox"/> Closed | | *STATUS_UNKNOWN | | vhost1165:LZ | | | | vhi | | | |
| | <input checked="" type="checkbox"/> Closed | | *STATIS UNKNOWN | | vhost1165:LZ | | | | vhi | | | |

Figure 4-26 Close alert status

Another situation alert will not be shown until you stop the ping process.

4.4 Running a CPU monitoring test for the Workload Governance and Management scenario

As part of Workload Governance and Management scenario, CPU monitoring is used to start the provisioning at the SmartCloud Environment. The concept is to detect high CPU usage and automatically respond to it by provisioning another virtual machine to handle the load. This section describes how to configure CPU monitoring to detect high usage.

First create the situation by selecting **System CPU Percent** and choosing the threshold value to generate the alert. In this example, the threshold is **higher than 50**. The situation is shown in Figure 4-27.

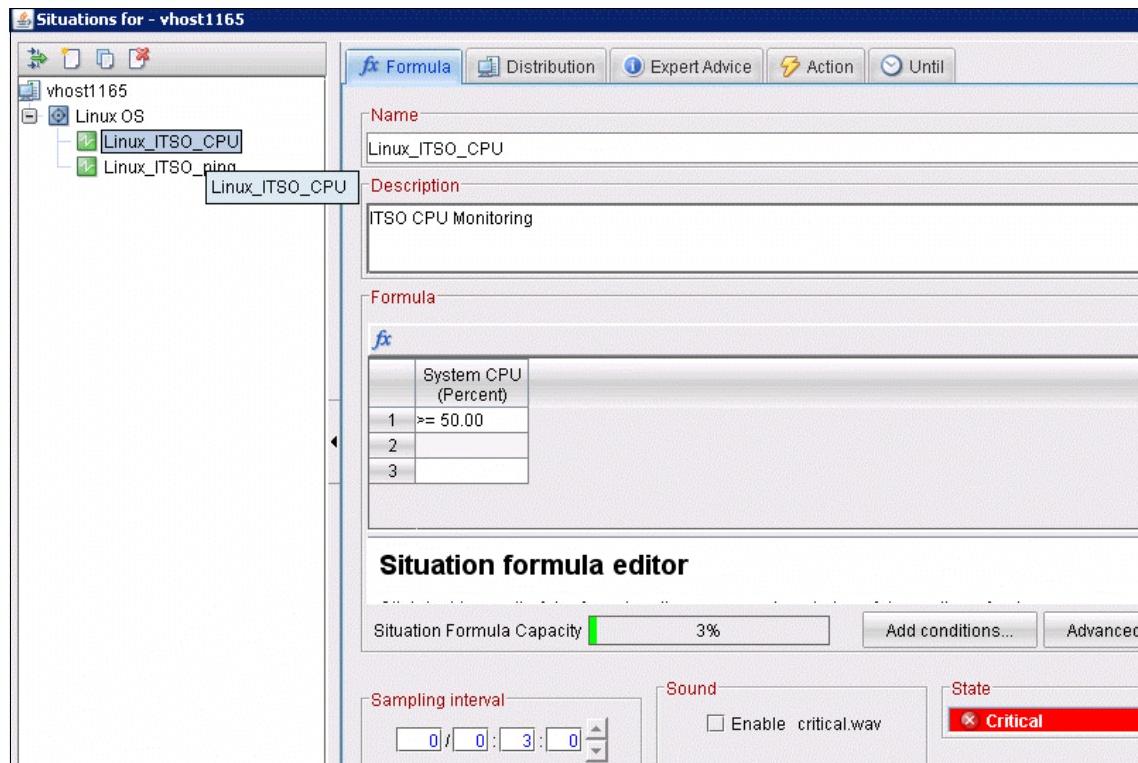


Figure 4-27 CPU monitoring situation

To test this situation, you can create a stress test on the Linux agent that will consume more than 50 percent of the available CPU. Alternatively, just to test the situation, you can monitor for a much lower threshold. After verifying the situation, raise the threshold to the desired level.

When a situation is triggered, IBM Tivoli Monitoring can be configured to respond automatically. In a production system, you would not react to a single spike in usage as shown in this example. IBM Tivoli Monitoring can be configured to react to trends rather than to transient behavior.

When high CPU usage is detected, the situation can be forward to Omnibus for further action. See Chapter 5, “Provisioning scenario” on page 93 for information about how to configure the response to the alert.

4.5 Lessons learned

This section discusses various lessons learned by the authors while setting up this scenario:

- ▶ After installing the extension plug-in, reload the browser.
- ▶ The log files are useful for diagnostics:
 - a. To access log files from the WebSphere Cast Iron integration appliance using the Web Management Console, select **System > Commands > Download postmortem** (Figure 4-28).
 - b. Select **Go** and then **Download**.

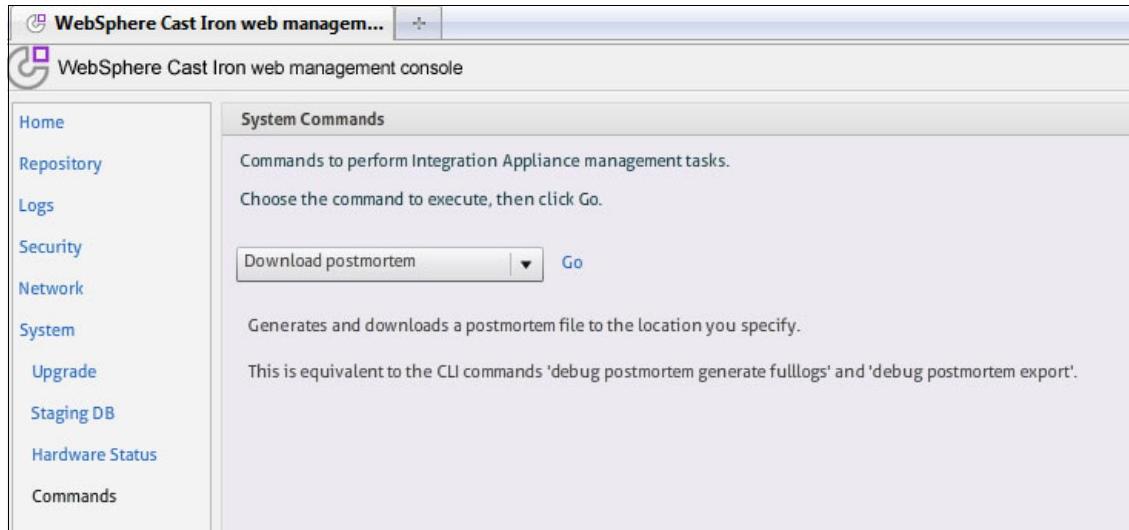


Figure 4-28 Command to download log files from integration appliance

- c. Save the file (Figure 4-29 on page 90).

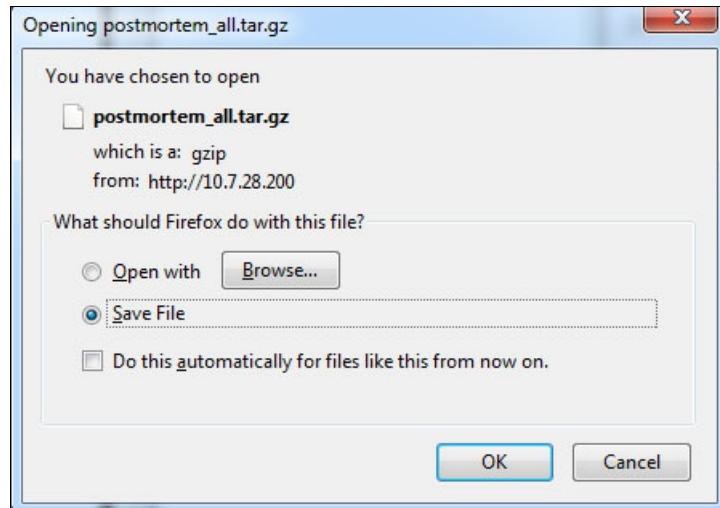


Figure 4-29 All log files from WebSphere Cast Iron integration appliance

d. The resulting file is `postmortem_all.tar.gz`. Uncompress this file.

The `postmortem` directory has three subdirectories:

- `usr`:
- `var`:
- `wip`: This directory has installation log files for the plug-ins.

The `postmortem\var\log` directory includes `wmc.out`, which can be helpful in finding connection problems (Figure 4-30).

```
integTEMSPortStatus value is :not connected
integWHPPortStatus value is :not connected
:10005    integDownStreamStatus value is :SYN_SENT
```

Figure 4-30 Example diagnostic information from `wmc.out` file

The `postmortem\wip\omniapps\plugins\install\ibm` directory has the following files and folders that you might find useful:

- `M9`: monitoring plug-in
 - `postmortem\wip\omniapps\plugins\install\ibm\M9\7.2.1.0\runtime\ITM\gateway\gateway.xml`
- `provisioningplugin`: provisioning plug-in
- `tdi`: Tivoli Directory Integrator

- ▶ On the WebSphere Cast Iron integration appliance, use the **netspect** command to diagnose problems, as shown in Figure 4-31.

```
ci-192-168-1-128/Standalone> netspect ping host 170.225.100.145
Pinging host 170.225.100.145 ...
PING 170.225.100.145 (170.225.100.145) 56(84) bytes of data.
64 bytes from 170.225.100.145: icmp_seq=1 ttl=128 time=71.9 ms
64 bytes from 170.225.100.145: icmp_seq=2 ttl=128 time=73.0 ms
64 bytes from 170.225.100.145: icmp_seq=3 ttl=128 time=71.4 ms
64 bytes from 170.225.100.145: icmp_seq=4 ttl=128 time=70.9 ms

--- 170.225.100.145 ping statistics ---
4 packets transmitted, 4 received, 0% packet loss, time 3015ms
rtt min/avg/max/mdev = 70.904/71.837/73.090/0.809 ms
```

Figure 4-31 Using the netspect command

- ▶ On the virtual machine, the **netstat -an** command helps identify port issues.

4.6 Preferred practices

This section discusses preferred practices for the monitoring plug-in.

- ▶ When choosing a name for a new integration, choose a unique and descriptive name.
- ▶ Identify necessary ports and make sure they are not being blocked by firewalls. Document which ports you are using, especially if you choose to change default values. Identify any network address translation (NAT) issues which impacts the addressing and plan for them.
- ▶ Identify any dependencies between the monitoring and other scenarios. For example, if provisioning new virtual machines, what are the monitored conditions that are to be reported to trigger provisioning? The criterion of 50 percent CPU utilization in this example is probably not adequate for a production system.
- ▶ Save and reuse virtual machines that have been properly configured.
- ▶ If you are planning to run many agents simultaneously, be aware that heartbeat and timeout values can impact performance. In the `lz.ini` file, the `CITIRA_HEARTBEAT` and `CTIRA_RECONNECT_WAIT` values may need to be tuned in systems with many agents.

4.7 Summary

In this chapter, the monitoring plug-in was installed and configured on the WebSphere Cast Iron integration appliance. A gateway agent was deployed to facilitate communication between the monitoring server on premises and the virtual machine operating system agents in the cloud.

A situation was created to verify that virtual machines in the cloud can be monitored on the same TEMS as local hosts.

The next step is to create a monitor to detect peak usage so corrective action can be taken.



Provisioning scenario

In this chapter we detail the Provisioning scenario of the fictitious, multinational retail company, IBM Redbooks Stores, which recently made an acquisition. As a result of this acquisition, the company needs to complete the following tasks:

- ▶ Integrate the IT systems of the acquired company with its own IT infrastructure
- ▶ Add new features into its existing applications to allow a common set of applications for both companies

In this scenario, to meet the first requirement, the Chief Information Officer (CIO) had to decide whether to integrate the new IT infrastructure in the company's data center or to deploy it off premise. The IBM Redbooks Stores data center was full and unable to support additional systems. Because creating new space takes time and is expensive, the CIO decided to deploy new systems using SmartCloud Enterprise, which was already in use for other workloads.

To meet the second requirement, the CIO told the development team to make changes with existing applications. To ensure that new features will not impact existing applications, the project manager required new environments for development, functional, and performance tests. The CIO decided to provide these temporary environments using SmartCloud Enterprise to avoid new capital expenses.

IBM Redbooks Stores currently has a private cloud managed platform on premise, based on IBM Service Delivery Management. Users are trained to use

the IBM Service Delivery Management administration portal, and the CIO does not want to add a new user experience to provision and release services through SmartCloud Enterprise. The CIO asked IBM to extend the existing capabilities of the portal for requesting a cloud service either on premise or through SmartCloud Enterprise. IBM agreed to add new services to the service catalog and provide these services through a unique portal.

In this chapter we discuss the following topics:

- ▶ Provisioning scenario overview
- ▶ Installing and configuring the Provisioning scenario
- ▶ Scenario usage
- ▶ Lessons learned
- ▶ Summary

5.1 Provisioning scenario overview

The Provisioning scenario provides the user with a single interface that has the ability to provision a cloud service either on premise or in a public cloud. So, the user has only one user experience and learning curve for all places where services are deployed. The Provisioning scenario uses extensions to Tivoli Service Automation Manager and provides the ability to provision servers within IBM Smart Cloud. Figure 5-1 illustrates an overview of this scenario.

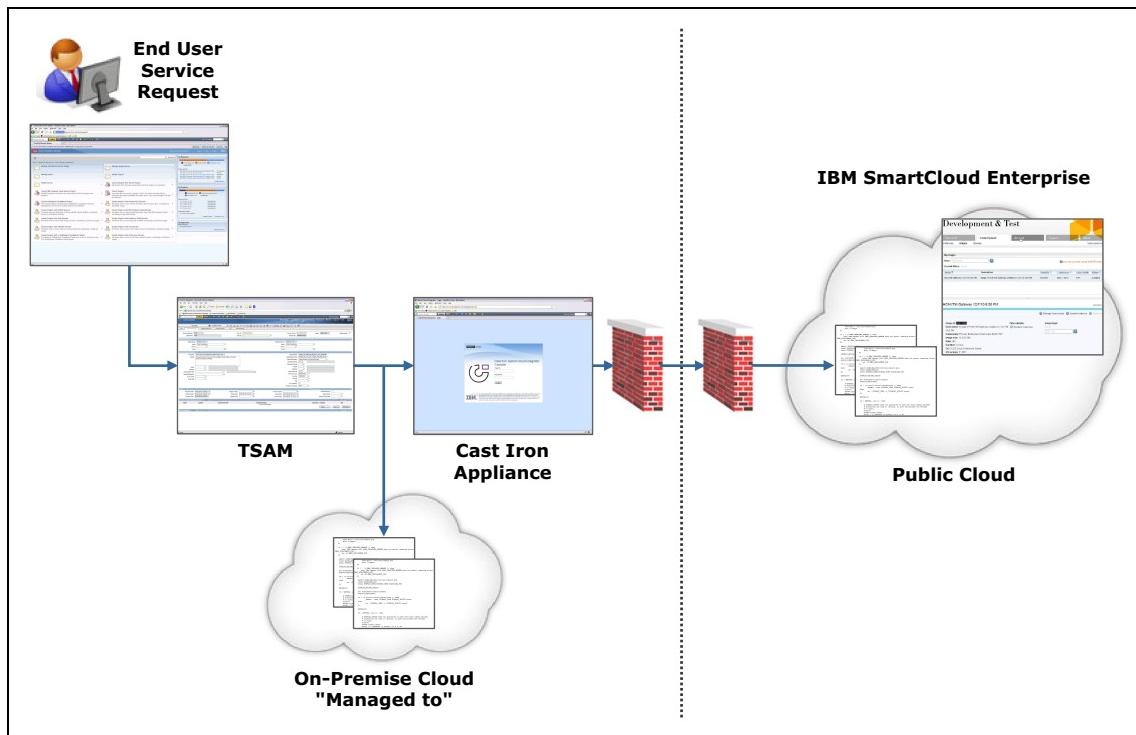


Figure 5-1 Provisioning scenario at glance

In the cloud management environment (Tivoli Service Automation Manager), two new service definitions are added. Management plans are then implemented to trigger provisioning and decommissioning.

The following new offerings are available:

- ▶ Create Project with IBM Compute Cloud Servers
- ▶ Add IBM Compute Cloud Server

In addition, the following existing offerings now support IBM SmartCloud:

- ▶ Remove Server
- ▶ Cancel Project

5.2 Installing and configuring the Provisioning scenario

At a high level, you need to complete the following steps to set up the Provisioning scenario:

- ▶ Installing the Service Management Extensions for Hybrid Cloud
- ▶ Creating the integration
- ▶ Installing the Tivoli Service Automation Manager Hybrid Cloud Extension

Important: The process described here covers only the installation steps that are relevant for the integration of the Provisioning scenario. It does not cover the full installation of the products.

5.2.1 Installing the Service Management Extensions for Hybrid Cloud

The first step is to install the Service Management Extensions for Hybrid Cloud. Refer to 2.8, “Installing the Service Management Extensions for Hybrid Cloud” on page 29 for installation instructions.

5.2.2 Creating the integration

To create an integration in the Cast Iron appliance that is used during the provisioning to public clouds, follow these steps:

1. Log in to the WebSphere Cast Iron Design Console using an address similar to the following, where *CIname* is the host name of the Cast Iron appliance or its IP address:

`http://CIname/express`

Use the following user name and password to log in:

Username: admin
Password: !n0r1t50C

Tip: If you have changed the password, use the new password that you created.

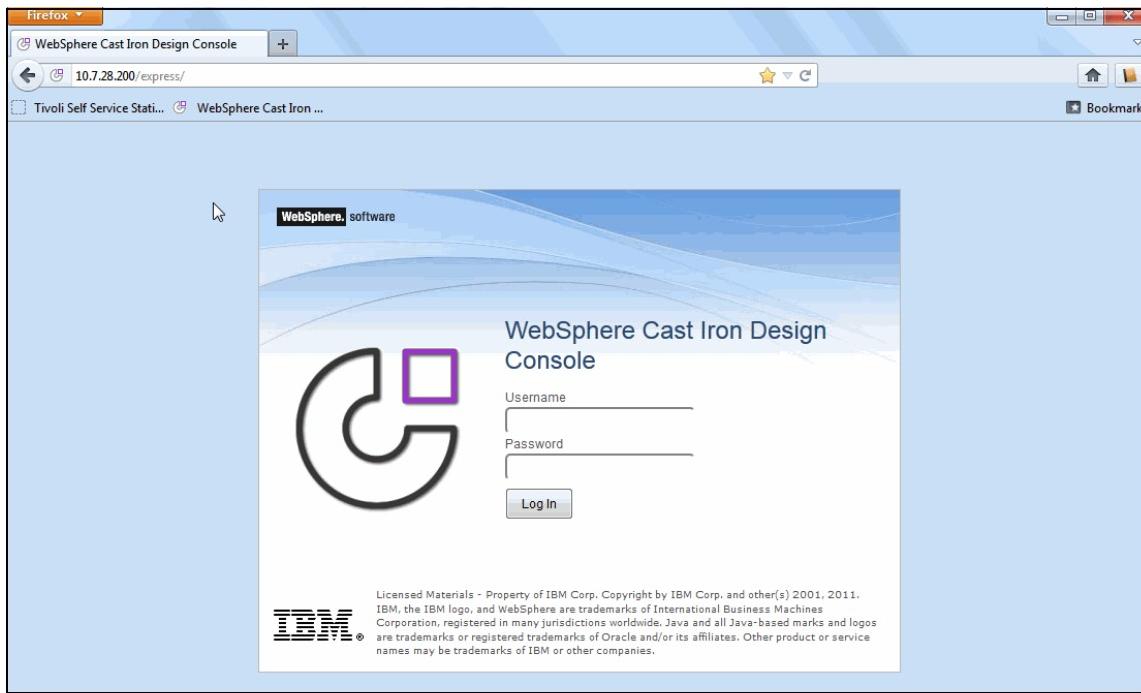


Figure 5-2 WebSphere Cast Iron Design Console login

2. Select **Source and Target**. Then complete these steps (Figure 5-3 on page 98):
 - a. Click **Integrations** in the toolbar to show the Create panel.
 - b. Select **TSAM** as the Source.
 - c. Select **IBM Compute Cloud** as the Target.
 - d. Click **New Blank Integration**.

Important: Do not use the Existing Connection drop-down list. Using this drop-down list will lead to unusual behavior of the appliance.

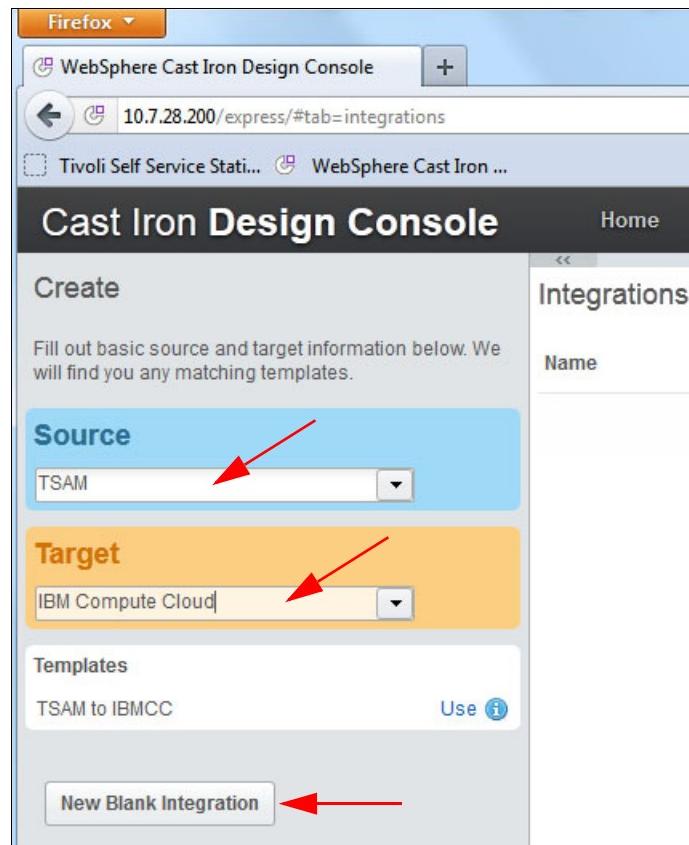


Figure 5-3 Select Source and Target

3. Save the integration.

4. After the integration is defined, name each endpoint (source and target) and save it (Figure 5-4).

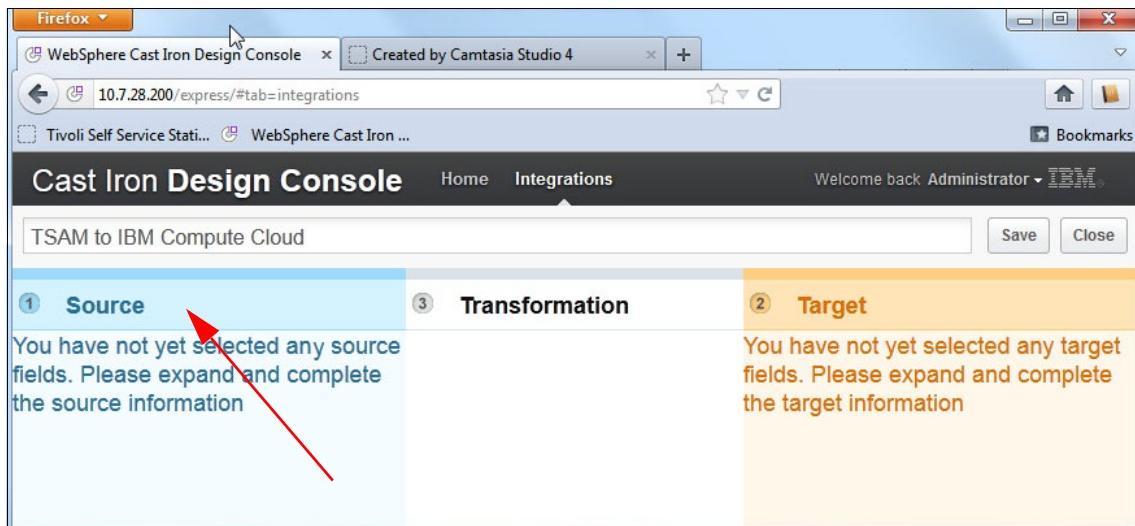


Figure 5-4 Finalize the integration

5. Click the **Source** tab to open the detail panel. Complete the details, and then click **Connect** (Figure 5-5).

The screenshot shows the 'Cast Iron Design Console' interface. The title bar says 'Cast Iron Design Console'. Below it, a header bar has 'TSAM to IBM Compute Cloud' on the left and 'Operations' on the right. A vertical sidebar on the right lists 'Available Connections' with 'IBMCCRAL' highlighted. The main area is titled '① Source'. It contains two tabs: 'Operations' (selected) and 'Available Connections' (disabled). Under 'Operations', there's a 'Connection Type' dropdown set to 'TSAM' and an 'Existing Connections' dropdown set to 'IBMCCRAL'. A '▼ Details' section contains fields for 'Name' (set to 'IBMCCRAL') and 'TSAM Description'. At the bottom is a 'Connect' button with a cursor pointing at it.

Figure 5-5 Complete the source details

6. Click the **Target** tab to open the detail panel. Complete the details, and then click **Connect** as shown in Figure 5-6.

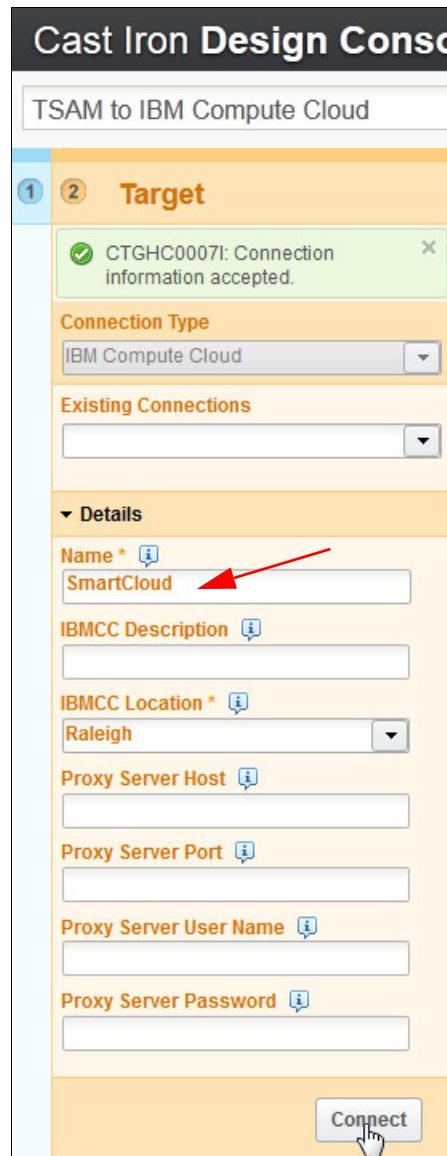


Figure 5-6 Complete the Target details

7. Click **Save**. A message indicates that the save was successful (Figure 5-7).

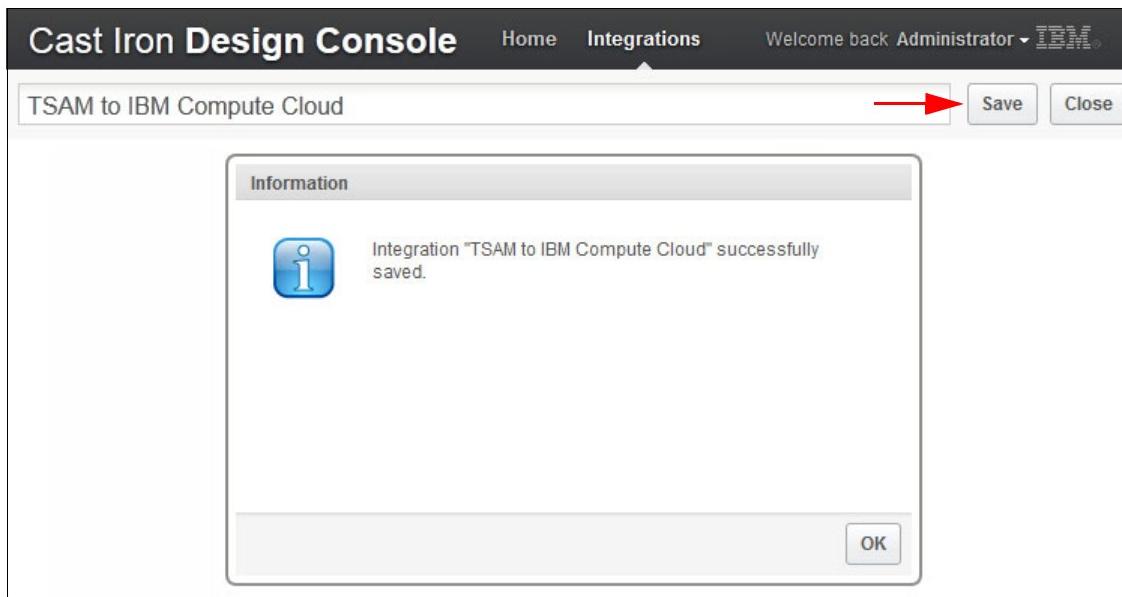


Figure 5-7 Save the integration

8. Click **Integrations**. The new integration displays in the console window in a stopped state, as shown in Figure 5-8. To start the integration, click the play icon. To stop it, click the stop icon. To delete it, click the delete icon.

| Name | Updated | Updated by | Status |
|---------------------------|-----------------------------------|------------|-------------|
| TSAM to IBM Compute Cloud | February 28, 2012 11:56:31 PM GMT | admin | Stopped ► ✖ |

Figure 5-8 Start the Integration

5.2.3 Installing the Tivoli Service Automation Manager Hybrid Cloud Extension

The provisioning of servers in a public cloud is managed by Tivoli Service Automation Manager and triggered by Tivoli Provisioning Manager workflows.

The installation of the service automation solution for hybrid cloud is a multistep process, as explained here:

1. Add custom Java code to the Tivoli Service Automation Manager to implement the operational workflow used by the solution.
2. Add a Dojo GUI to extend the simple Tivoli Service Request Manager user interface.
3. Install a new Tivoli Provisioning Manager driver that contains the Tivoli Provisioning Manager workflows. The Tivoli Provisioning Manager workflow implements the server provisioning and decommissioning, and the image discovery processes.
4. Add new service definitions. Add the associated actions, operational workflows, job plans, and topology nodes.
5. Create offerings and add them to the offerings catalogs. Add the associated classifications.
6. Create escalation definitions to implement workload governance and management of servers that run in a hybrid cloud.

Downloading the Hybrid Cloud extension software

To download the Hybrid Cloud extension software, follow these steps:

1. Log on to the Integrated Service Management Library, and search for Hybrid Cloud.
<https://www-304.ibm.com/software/brandcatalog/ismlibrary/>
2. When you locate the Service Management Extensions for Hybrid Cloud offering, click **Download**, as shown in Figure 5-9.

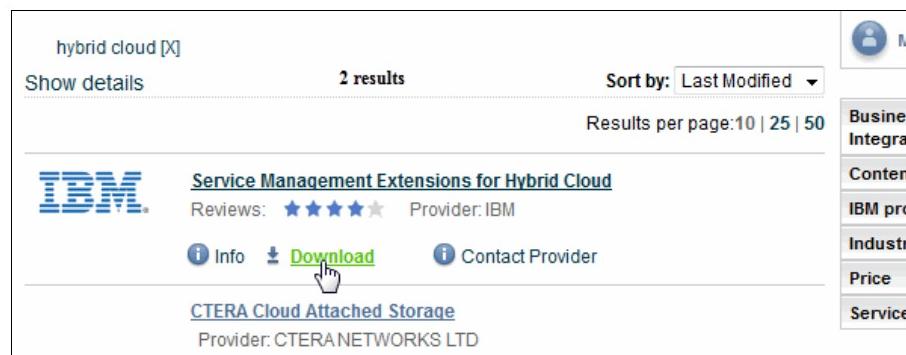


Figure 5-9 Download the Hybrid Cloud extension

3. Select the **Service Management Extensions for Hybrid Cloud 7.2.1**, as shown in Figure 5-10.

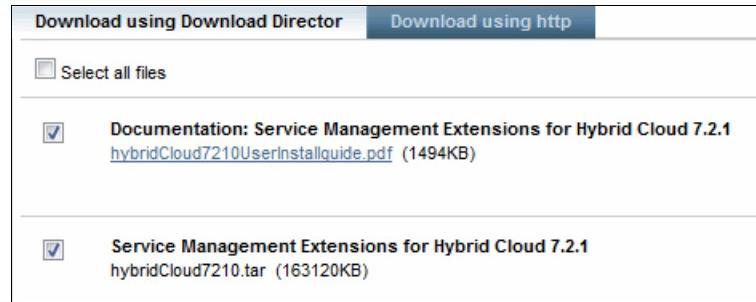


Figure 5-10 Extensions selection

To download the Tivoli Service Automation Manager Hybrid Cloud extension software:

1. Ensure that the Tivoli Service Automation Manager V7.2.1.3 server is running.
2. Log on to the Tivoli Service Automation Manager host as the root user.
3. Transfer the hciExtension.tar file from the original media to the /tmp directory. Then, change to the /tmp directory:

```
# cd /tmp
```
4. Extract the contents of the hciExtension.tar file to create an hci directory that contains all the necessary files.

```
# tar xvf hciExtension.tar
```
5. Check that the files and folders shown in Figure 5-11 are available.

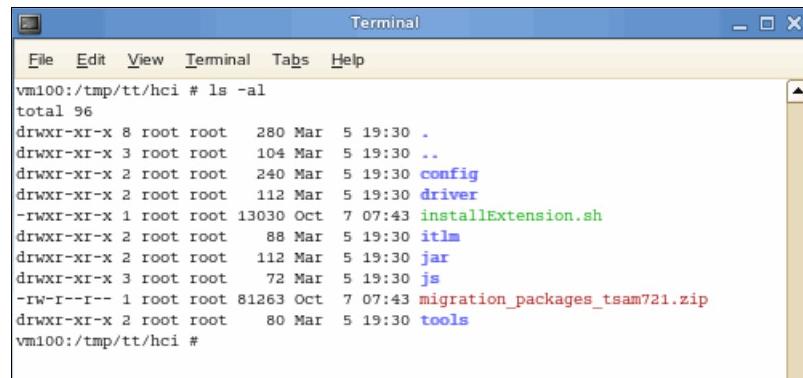


Figure 5-11 Files unzipped

Installing the Tivoli Provisioning Manager driver

To install the Tivoli Provisioning Manager driver, complete the following steps:

1. Log on to the Tivoli Service Automation Manager host as the root user.

2. Change directory to the /tmp/hci directory:

```
# cd /tmp/hci
```

3. Run the installExtension.sh script:

```
# ./installExtension.sh
```

4. The installExtension.sh script assumes the following default values:

| | |
|----------------|---------------------|
| eWASLogin | wasadmin |
| eWasPasswd | password |
| eWebSphereHome | /opt/IBM/WebSphere |
| tmpHome | /opt/IBM/Tivoli/tmp |

If the values are different in your system, use the appropriate command line arguments to pass the correct values. Use the `./installExtension.sh -h` command to display the usage page. For example:

```
# ./installExtension.sh -l wasadmin -p mypassword -w /app/WebSphere  
-t /app/tpm
```

A message displays if the installation completes successfully (Figure 5-12).

```
MXServer is started
Start TPM engines
    ...waiting
    ...waiting
The lightweight runtime has started successfully.
TPM startup completed.

Check the following logs for startup errors:
/usr/ibm/tivoli/common/COP/logs/tio_start.log
/usr/ibm/tivoli/common/COP/logs/tio_start_service.log
/opt/IBM/tivoli/tpm/lwi/logs
/opt/IBM/WebSphere/AppServer/profiles/ctgDmgr01/logs/dmgr
/opt/IBM/WebSphere/AppServer/profiles/ctgAppSrv01/logs/nodeagent
/opt/IBM/WebSphere/AppServer/profiles/ctgAppSrv01/logs/MXServer
/opt/IBM/WebSphere/AppServer/profiles/casprofile/logs/server1
Info: Installation completed successfully. ←
vm100:/tmp/hci # 1
```

Figure 5-12 Extension Installation complete

Importing the migration packages

Next, you need to import the migration packages for Tivoli Service Automation Manager 7.2.1 Fix Pack 3.

Prerequisite steps

Before you import the migration packages, complete the following steps:

1. Log on to the server where you run your web browser. Then, extract the following files from the `migration_packages_tsam721.zip` file from the `/tmp/hci` directory (Figure 5-13):
 - `HCIArtifact-Base.zip`
 - `HCIArtifact-Additional.zip`

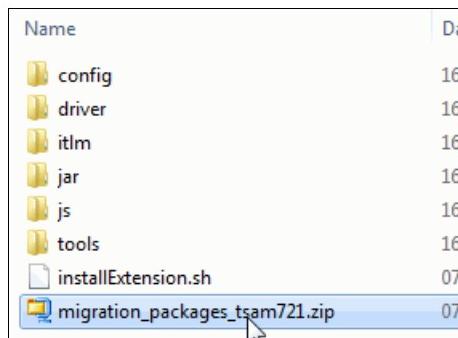


Figure 5-13 Extract the migration package

2. Log on to IBM Maximo® Start Center as described in 5.3.6, “Logging in to the Maximo Start Center” on page 150.
3. Check your version of Tivoli Service Automation Manager by clicking **Help > System Information**. The version displays as shown in Figure 5-14 on page 107.

Important: You must use Tivoli Service Automation Manager V7.2.1.3.

| System Information | | ? |
|------------------------------|--|---|
| App Server | IBM WebSphere Application Server 6.1 | |
| User Name | MAXADMIN | |
| Version | Service Desk Integration MEA 7.2.0.0 Build 20100419D DB Build V7200-03 TPM PMP 7.2.0.0-IF00001 Build U3133_40-20100810 DB Build V7201-05 SRM Service Request Management 7.2.0.1 Build 20100316D2 DB Build V7201-01 RDP PMP 7.2.1.3 Build 20101217D DB Build V724-06 IBM Tivoli Change Management Content 7.2.0.0 Build 20091111D DB Build V7200-14 IMBot SRM Service Desk 7.2.0.0 Build 20100419D DB Build V7200-03 SRM Service Desk Content - Best Practices 7.2.0.1 Build 20100316D2 DB Build V7201-01 Service Automation Manager - CCMDB Integration PMP 7.2.1.2 Build 20101217D DB Build V723-03 SRM Service Desk Content - Classification 7.2.0.0 Build 20100419D DB Build V7200-05 Service Automation Manager PMP 7.2.1.3 Build 20101217D DB Build V724-05 IBM Tivoli Common Process Components 7.2.0.0 Build 20100109D DB Build V7201-02 Service Automation Manager Configuration PMP 7.2.1.3 Build 20101217D DB Build V724-04 SRM Service Desk Content - Best Practice Users 7.2.0.0 Build 20100419D DB Build V7200-05 SRM Service Desk 7.2.0.1 Build 20100316D2 DB Build V7201-04 SRM Screen Capturer 7.2.0.0 Build 20100419D DB Build V7200-02 SRM Problem Management 7.2.0.1 Build 20100316D2 DB Build V7201-01 SRM Solution 7.2.0.1 Build 20100316D2 DB Build V7201-01 SRM Incident Management 7.2.0.1 Build 20100316D2 DB Build V7201-01 SRM Service Catalog Base 7.2.0.1 Build 20100316D DB Build V7201-05 SRM SLA Hold 7.2.0.1 Build 20100316D2 DB Build V7201-03 SRM Search 7.2.0.1 Build 20100316D2 DB Build V7201-02 CI Reservation PMP 7.2.1.0 Build 20101217D DB Build V711e-31 SRM Survey Management 7.2.0.1 Build 20100316D2 DB Build V7201-01 IBM Tivoli UI Widgets 7.2.0.01 Build 201002251643 DB Build V7201-02 IBM Tivoli Change Management 7.2.0.01 Build 20100215D2 DB Build V7201-02 ServiceProvider 7.1.1.1 Build BUILD DB Build V7121-03 HFDB Build HF7121-01 Service Automation Manager WAS PMP 7.2.1.0 Build 20100702D DB Build V721-05 SRM Service Catalog Content 7.2.0.1 Build 20100316D DB Build V7201-06 Base Services 7.1.1.6-LA20100521-0608 Build 20091208-1415 DB Build V7116-173 HFDB Build HF7116-08 | |
| Server OS | Linux 2.6.16.60-0.54.5-smp | |
| License # | | |
| (c) Copyright IBM Corp. 2007 | |  |

Figure 5-14 Tivoli Service Automation Manager version

Updating the CLASSUSETWITH domain

To update the CLASSUSETWITH domain, complete the following steps:

1. Log on to Maximo Start Center.
2. Click **Go To > System Configuration > Platform Configuration > Domains**.

3. In the **Domain** field in the Filter area, specify CLASSUSEWITH and press Enter. An entry displays in the table.
4. Click **Edit Detail** icon to the left of the CLASSUSEWITH domain name. Then, in the **Value** field in the Filter area of the SYNONYM Domain window, specify PMSCCRSPEC and press **Enter**. If the PMSCCRSPEC entry exists, no further action is required.
5. If the PMSCCRSPEC entry does not exist, create the entry as follows:
 - a. Click **New Row**.
 - b. Then, in the **Internal Value** field, specify TKTEMPLATE, and in the **Value** field, specify PMSCCRSPEC. Refer to Figure 5-15.
 - c. In the **Description** field, specify Use with Specification display table.
 - d. Click **OK**. Click the **Save Domain** icon.

The screenshot shows the 'SYNONYM Domain' configuration window. At the top, there are fields for 'Domain' (CLASSUSEWITH), 'Classification Use With', 'Domain Type' (SYNONYM), and 'Length' (30). Below this is a table with two rows:

| | Internal Value | Value |
|--|----------------|------------|
| | | PMSCCRSPEC |

Below the table, a new row is being created with the following values:

| | | |
|--------------------------|------------|------------|
| <input type="checkbox"/> | TKTEMPLATE | PMSCCRSPEC |
|--------------------------|------------|------------|

At the bottom, a modal dialog box is open with the following fields:

- * Internal Value: TKTEMPLATE
- * Value: PMSCCRSPEC
- Description: use with Specification display table

A red box highlights the 'Description' field and its value.

Figure 5-15 CLASSUSEWITH domain

Importing the base migration package

To import the base migration package:

1. Log on to Maximo Start Center.
2. Click **Go To > System Configuration > Migration > Migration Manager**.

3. Click **Select Action > Upload Package**.

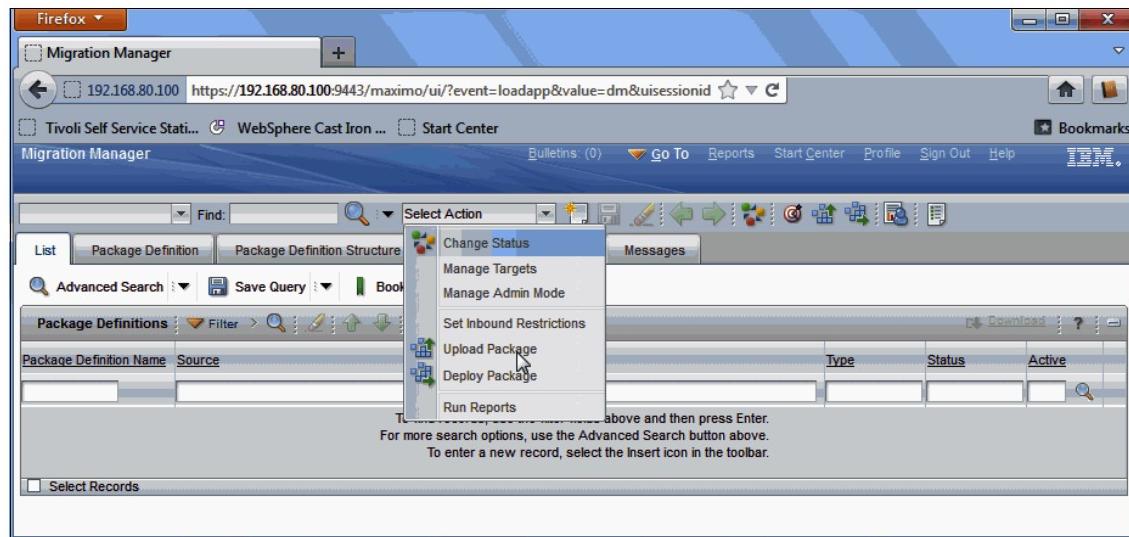


Figure 5-16 Upload package

4. In the Upload Package window, select the following file in the /tmp/hci directory:
`HCIArtifact-Base.zip`
5. To upload the file to the Tivoli Service Automation Manager server, click **OK**.
6. Click **Select Action > Deploy Package**.
7. In the Deploy Package dialog box, select the **Do you have a current backup** option. Then, click **Deploy** (Figure 5-17).

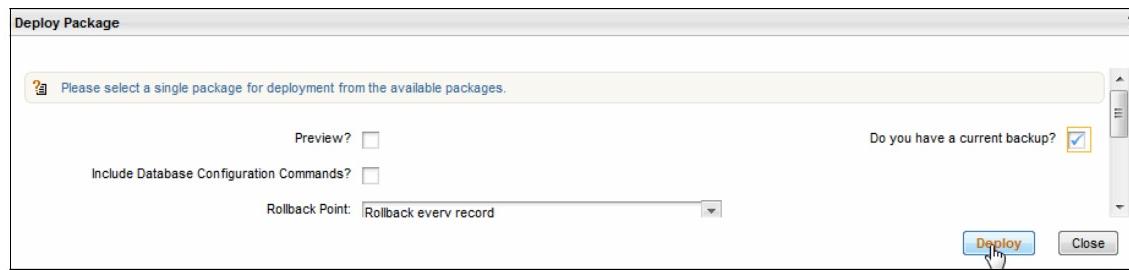


Figure 5-17 Deploy package

To confirm that the package deployed successfully, verify that you receive the messages shown in Figure 5-18 and Figure 5-19 on page 111.

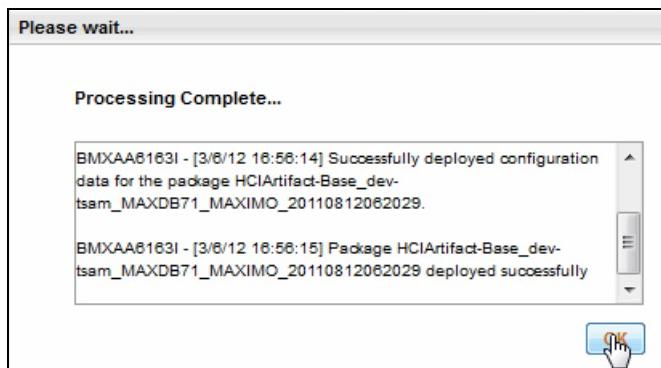


Figure 5-18 Successful deployment

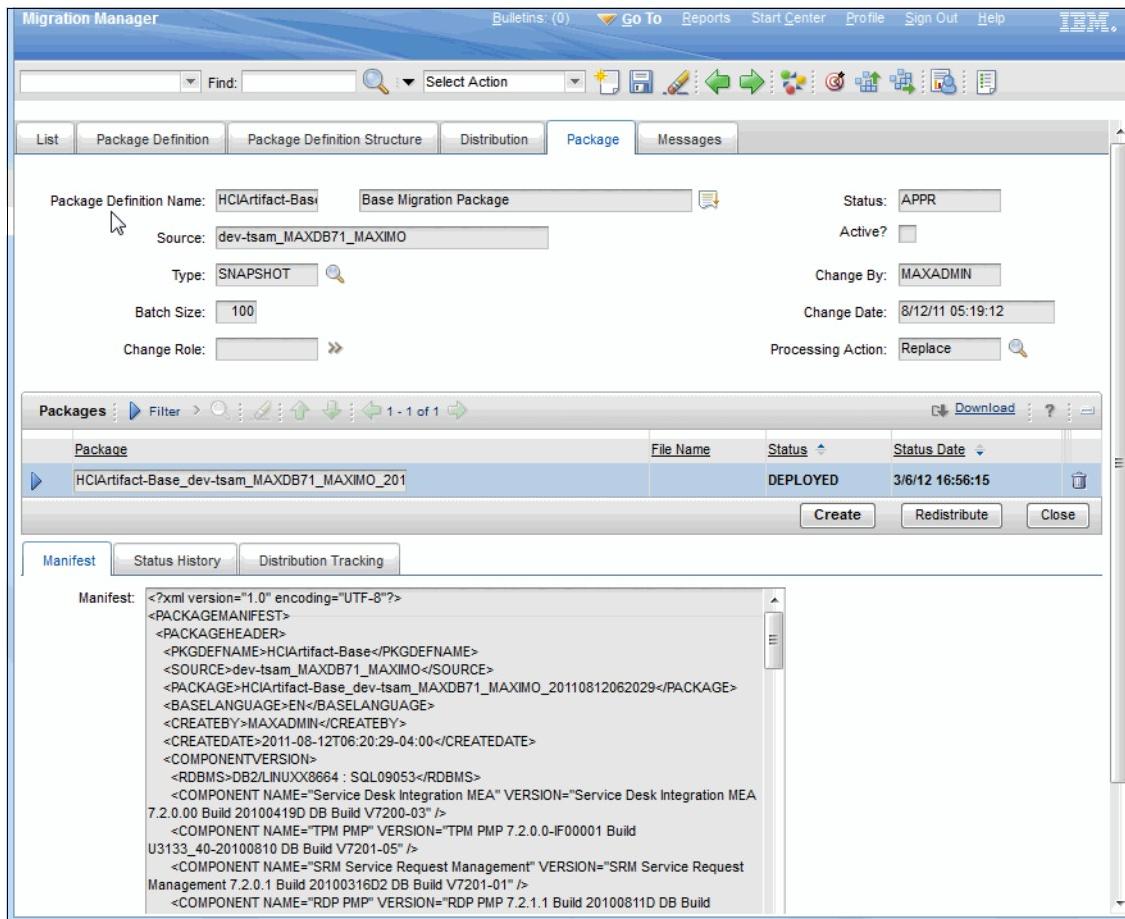


Figure 5-19 Package deployed

8. Click **Go To > Service Automation > Service Definitions**.

9. Verify that the service definitions list contains the GENICC service definition, as shown in Figure 5-20.

The screenshot shows a software application window titled "Service Definitions". At the top, there is a toolbar with various icons for search, selection, and file operations. Below the toolbar is a menu bar with tabs: "List", "Service Definition", "Management Plans", "Monitoring Information", "Service Topology", and "Service Catalog". Underneath the menu bar, there are search and filter options: "Advanced Search", "Save Query", and "Bookmarks". The main area is a table titled "Service Definitions" with the following columns: "Service Definition ID", "Name", "Revision", "Latest Revision", and "Description". There are two rows in the table. The first row has a yellow background and contains the value "ICC" in the "Service Definition ID" column. The second row, which is highlighted with a red border, contains the values "GENICC", "IBM Compute Cloud Virtual Server Provisioning", "1", and "IBM Compute Cloud Virtual Server Provisioning". A checkbox labeled "Select Records" is located at the bottom left of the table. The status bar at the bottom of the window displays the message "1 - 1 of 1".

| Service Definition ID | Name | Revision | Latest Revision | Description |
|-----------------------|---|----------|-------------------------------------|---|
| ICC | | | | |
| GENICC | IBM Compute Cloud Virtual Server Provisioning | 1 | <input checked="" type="checkbox"/> | IBM Compute Cloud Virtual Server Provisioning |

Figure 5-20 New service definition

10. Select the GENICC service definition and then clicking the **Management Plans** tab. Verify that the following management plans are listed (Figure 5-21 on page 113):

- ADDSERVERICC
- CANPROJECT
- DELSERVER
- NEWICC

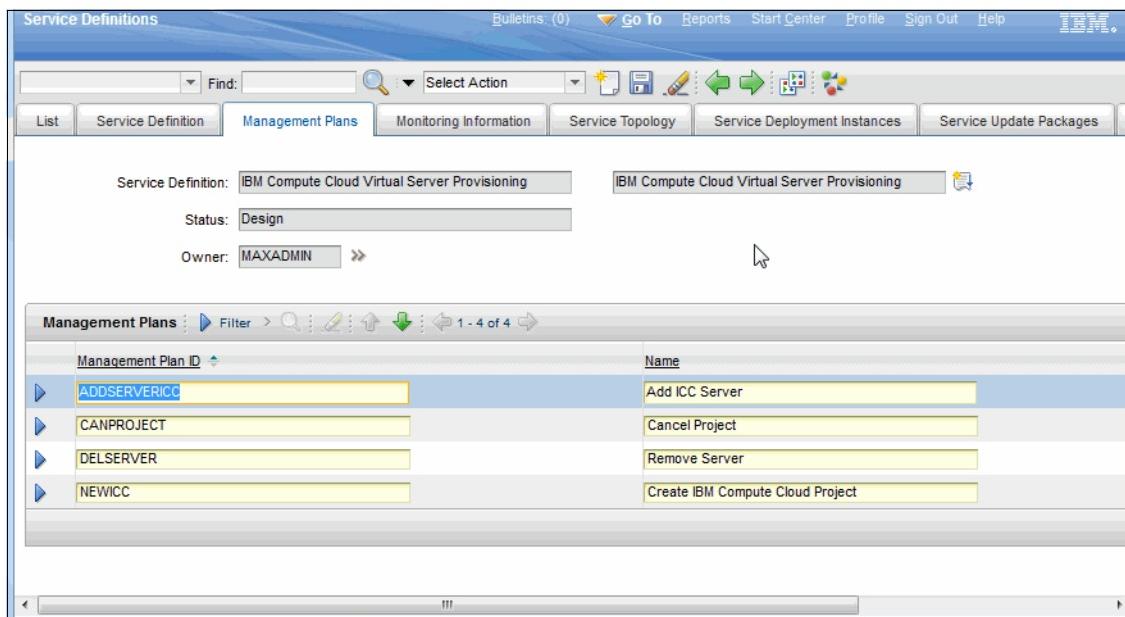


Figure 5-21 New management plan ICC

11. Click **Go To > Planning > Job Plans**. Verify that the job plans are listed as shown in Figure 5-22.

| Job Plans | | | | | | |
|-------------|--|----------|------------|--------|---------------|--------------|
| Job Plans | | | | | | |
| Job Plan | Description | Duration | Supervisor | Status | Template Type | Organization |
| ICC | | | | | | |
| PMRDPPDICC | Deploys IBM Compute Cloud Servers | 0:00 | | ACTIVE | ACTIVITY | |
| PMRDPPDSICC | Destroys Deployed IBM Compute Cloud Servers | 0:00 | | ACTIVE | ACTIVITY | |
| PMRDPPPICC | Get Properties of Deployed IBM Compute Cloud | 0:00 | | ACTIVE | ACTIVITY | |
| PMRDPPRICC | Personalizes virtual server nodes for ICC Deployment | 0:00 | | ACTIVE | ACTIVITY | |

Figure 5-22 Job plan ICC

Importing the additional migration packages

To import the HCIArtifact-Additional.zip migration package, follow these steps:

1. Log on to Maximo Start Center.
2. Click **Go To > System Configuration > Migration > Migration Manager**.
3. Click **Select Action > Upload Package**.
4. In the Upload Package window, select the following file in the /tmp directory:
HCIArtifact-Additional.zip
5. To upload the file to the Tivoli Service Automation Manager server, click **OK**.
6. Click **Select Action > Deploy Package**.
7. In the Deploy Package dialog box, select the **Do you have a current backup** option. Then, click **Deploy** as shown in Figure 5-23.

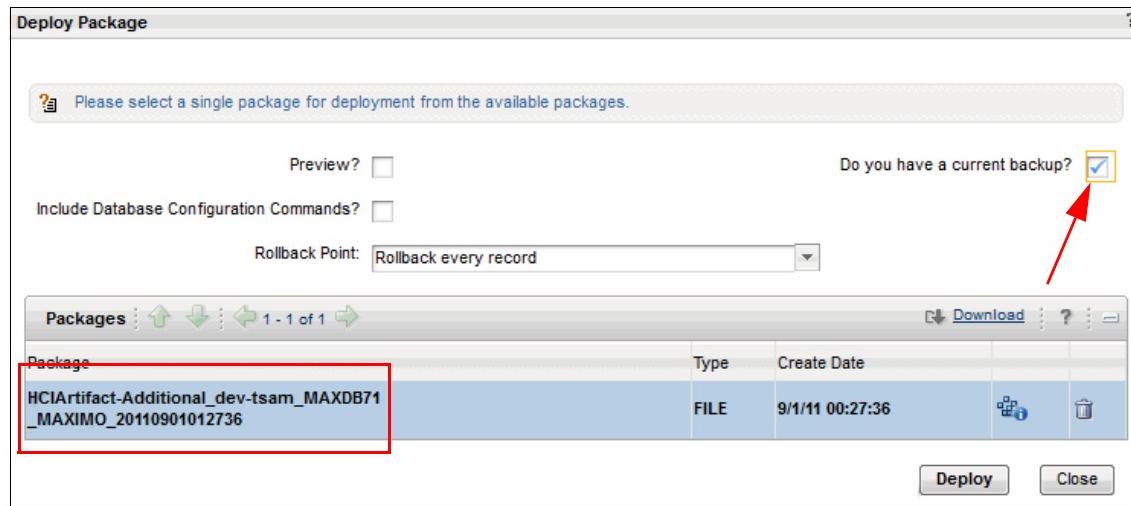


Figure 5-23 Deploy Additional Package

To confirm that the package deployed successfully, verify that you receive the message shown in Figure 5-24.

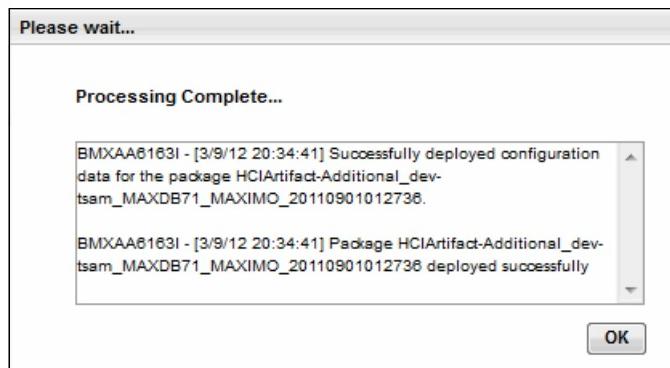


Figure 5-24 Successful additional deployment

Perform the following steps to verify the Additional migration package installation:

1. Click **Go To > Administration > Classifications**. Verify that the classifications list includes the classifications shown in Figure 5-25.

| Classifications | |
|---|---|
| Classification | Description |
| CLOUD | |
| PMRDP SR CREATE PROJECT PUBLIC CLOUD SERVER | Create Project with Servers in Public Cloud |
| PMRDP SR ADD PUBLIC CLOUD SERVER | Add Servers in Public Cloud |
| PMRDP SR CREATE PROJECT HYBRID CLOUD SERVER | Create Project with Servers in Hybrid Cloud |
| <input type="checkbox"/> Select Records | |

Figure 5-25 Classifications

2. Click **Go To > System Configuration > Platform Configuration > Actions**.
Verify that the actions list includes the actions shown in Figure 5-26.

The screenshot shows a software interface titled "Actions". At the top, there are search and filter fields labeled "Find:" and "Select Action". Below this is a toolbar with buttons for "Advanced Search", "Save Query", and "Bookmarks". The main area is a table titled "Actions" with columns for "Action" and "Description". The table contains the following data:

| Action | Description | Obj |
|----------------------|--|-----|
| ICC | | |
| <u>PMRDPPPICC</u> | ICC Prepare input for deployment | PM |
| <u>PMRDPPIRICC</u> | ICC Personalize virtual server node for Deployment | PM |
| <u>PMRDPICCP</u> | Scheduling extension for ICC Deployments | SR |
| <u>PMRDPAASVRICC</u> | Add ICC server to project | SR |
| <u>PMRDPSVRIICC</u> | Delete ICC Server | SR |

Figure 5-26 Actions

- Click **Go To > System Configuration > Platform Configuration > Workflow Designer**. Verify that the workflow list includes the workflow processes shown in Figure 5-27.

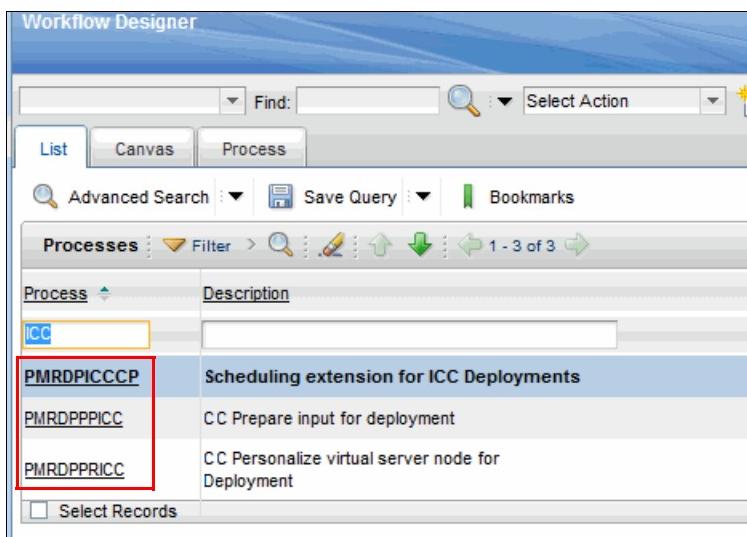


Figure 5-27 Workflows processes

- Click **Go To > Service Request Manager Catalog > Offerings**. Verify that the offerings list includes the offerings shown in Table 5-1.

Table 5-1 Offerings

| Offering | Description |
|--------------|---|
| PMRD_P_0751B | Create Project with IBM Compute Cloud Servers |
| PMRD_P_0752B | Add IBM Compute Cloud Server |
| PMRD_P_0802B | Create Project with Hybrid Servers |

Modifying the service definitions

Now you need to modify the service definitions that were installed as part of the Tivoli Service Automation Manager Hybrid Cloud Extension:

- Log on to Maximo Start Center.
- Click **Go to > Service Automation > Service Definitions**.
- In the Service Definition ID field in the Filter area, specify GENICC and press Enter. An entry displays in the table.
- Select the **GENICC** entry to open the service definition.

- Click the **Management Plans** tab and select **NEWICC** (Figure 5-28).

The screenshot shows the Service Definitions interface with the Management Plans tab selected. At the top, there are fields for Service Definition (IBM Compute Cloud Virtual Server Provisioning), Status (Design), and Owner (MAXADMIN). Below this, the Management Plans section displays a list of management plans:

| Management Plan ID | Name |
|--------------------|----------------------------------|
| ADDERVERICC | Add ICC Server |
| CANPROJECT | Cancel Project |
| DELSERVER | Remove Server |
| NEWICC | Create IBM Compute Cloud Project |

At the bottom, the Management Plan Details section shows the selected Management Plan ID as NEWICC.

Figure 5-28 GENICC Management Plan

6. In the Management Plan Details area, select the **initial** option (as shown in Figure 5-29). Then click **Save**.

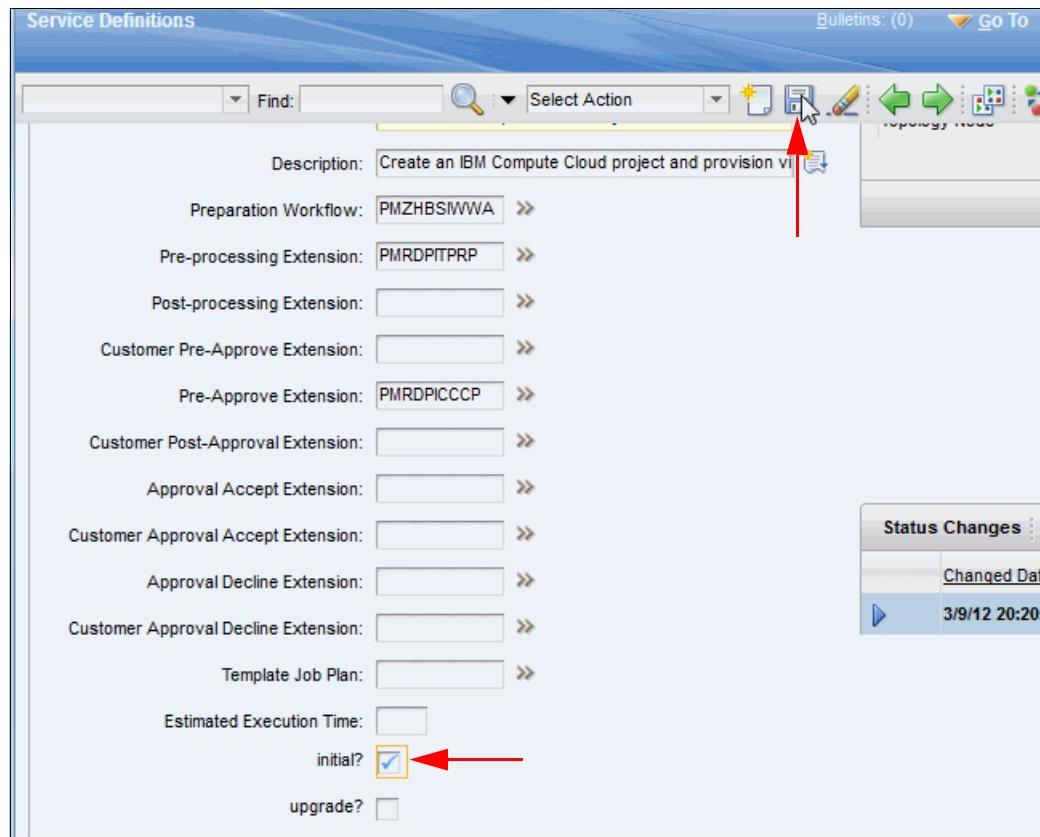


Figure 5-29 Initial option

Modifying the offerings

Next, update the offerings that were installed as part of the Tivoli Service Automation Manager Hybrid Cloud Extension, and add them to the relevant catalogs:

1. Log on to Maximo Start Center.
2. Click **Go to > Service Request Manager Catalog > Offerings**.
3. In the **Offering** field in the Filter area, specify PMRDP_0751B and press Enter.
4. Select the **PMRDP_0751B** entry to open the offering definition.
5. From the **Select Action** list above the tabs, select **Change Status**.

6. From the **New Status** list in the Change Status window, select **Pending** (Figure 5-30), and then click **OK**.

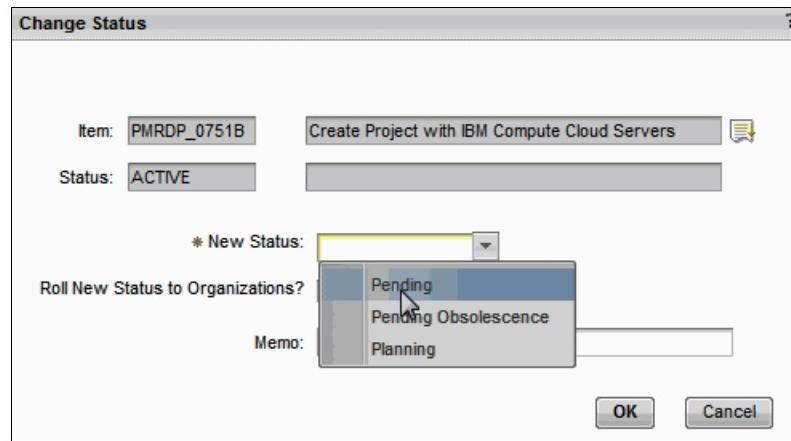


Figure 5-30 Change status

7. From the **Select Action** list above the tabs, select **Add Offering to catalog**, as shown in Figure 5-31.

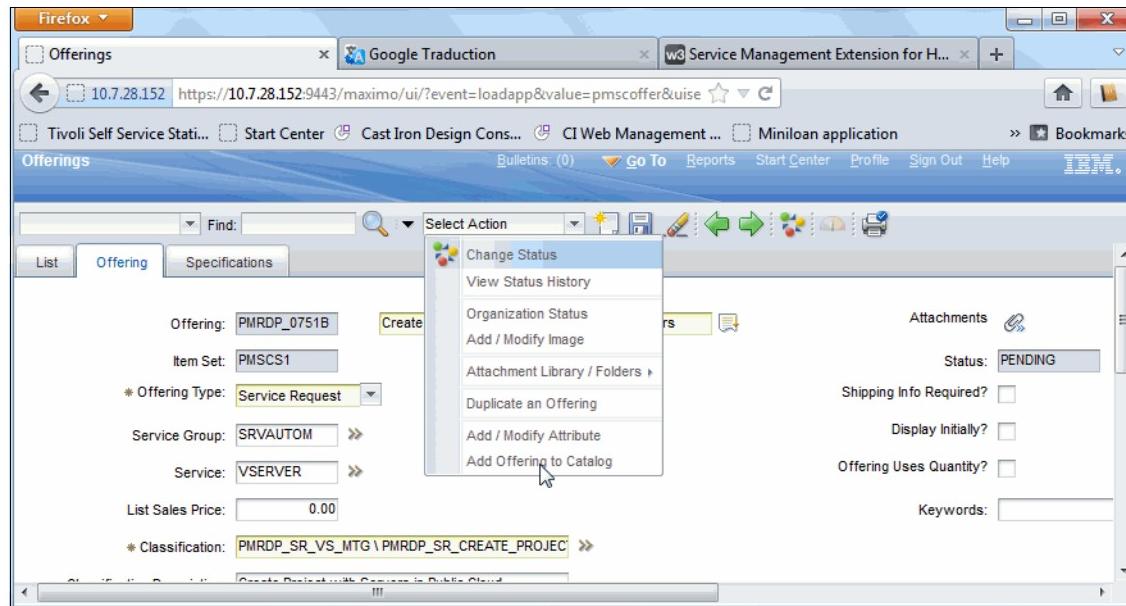


Figure 5-31 Add Offering to Catalog

- In the Add To Catalog window, click **Add to Catalog**. Then, in the empty row, click the **Detail Menu** icon (Figure 5-32).

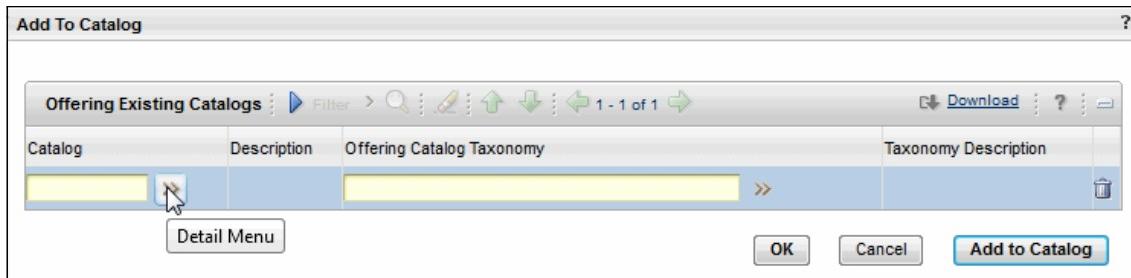


Figure 5-32 Detail Menu

- Click **Select a Value**. In the Select Value window, select **PMRDPCLOUDADMCAT_72**. Then, repeat steps 8 – 11 to add the offering to the PMRDPTEAMADMCAT_72 catalog. Then, click **OK**. The offering is added to the two catalogs.
- In the Service Request Information pane in the Offerings window, select the **Automatic Line Manager Approval** option (Figure 5-33).

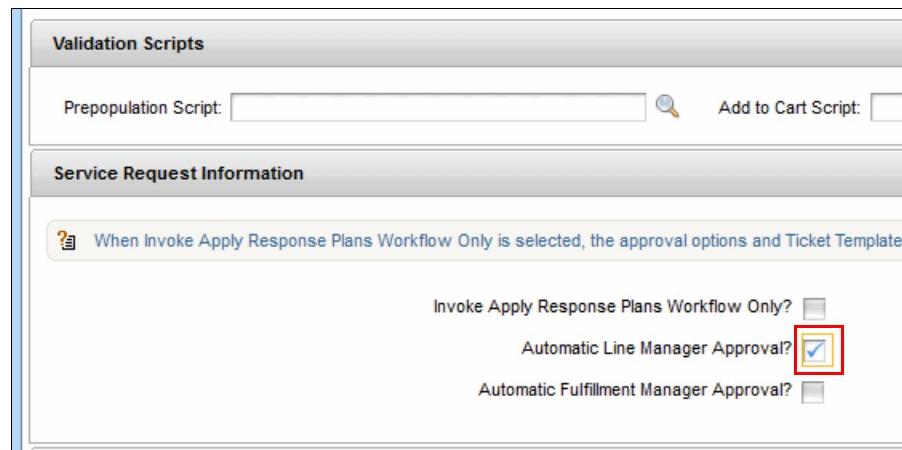


Figure 5-33 Automatic Line Manager Approval

- From the **Select Action** list above the tabs, select **Change Status**. Then, change the status of the offering to **Active**, and click **OK**. The offering is configured.
- Repeat steps 2 – 17 to configure the following offerings:
 - PMRDP_0752B: Add IBM Compute Cloud Server
 - PMRDP_0802B: Create Project with Hybrid Servers

Activating operational workflows

Now enable and activate the operational workflows that were installed as part of the Tivoli Service Automation Manager Hybrid Cloud Extension:

1. Log on to Maximo Start Center.
2. Click **Go To > System Configuration > Platform Configuration > Workflow Designer**.
3. In the **Process** field in the Filter area, specify ICC and press Enter.
4. Select the **PMRDPICCCP** entry to open the workflow definition.
5. From the **Select Action** list above the tabs (Figure 5-34), select **Enable Process**.

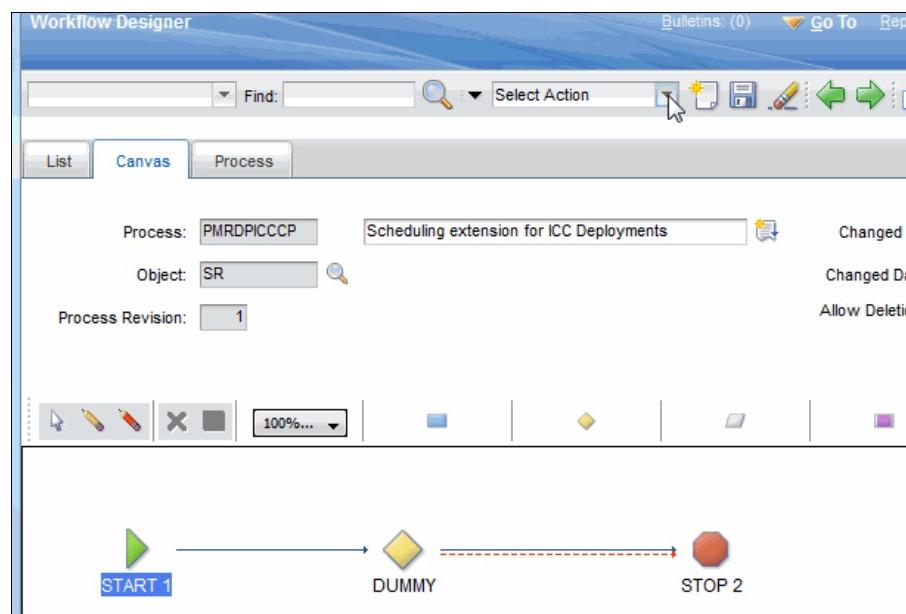


Figure 5-34 Enable Process

6. Click the **Canvas** tab, and verify that the **Enabled** check box selected.
7. From the **Select Action** list above the tabs, select **Activate Process**.
8. In the **Canvas** tab, verify that the **Active** option is selected.
9. Click the **List** tab.
10. Repeat steps 4 – 8 for the following workflows:
 - PMRDPPPICC
 - PMRDPPRICC

5.2.4 Configure Tivoli Service Automation Manager for a hybrid cloud

You need to update Tivoli Service Automation Manager to enable communication with the IBM SmartCloud Enterprise cloud through the WebSphere Cast Iron Appliance.

The provisioning of servers in a public cloud is managed by Tivoli Service Automation Manager and triggered by Tivoli Provisioning Manager workflows. The WebSphere Cast Iron Appliance provides the interface between the enterprise and the public cloud.

Configuring the service automation solution for hybrid cloud is a multistep process:

- ▶ Creating provisioning servers

The provisioning servers store the following information:

- Information about the environment, such as the WebSphere Cast Iron Appliance
- Information about the public cloud, such as user names and passwords and access keys.

- ▶ Creating and enabling cloud pools

A cloud pool is a representation of a cloud, either a private cloud such as VMware, or a public cloud such as IBM SmartCloud Enterprise

- ▶ Creating and registering images

The Tivoli Service Automation Manager images are the representation of real images in the public cloud and are used for provisioning

Creating provisioning servers

The provisioning servers are defined in the `DCM_objects.xml` file. A sample `DCM_objects.xml` file is provided in the `/tmp/hci/config` directory when you extract the `hciExtension.tar` file.

Locate on same system: The `DCM_objects.xml` file must be located on the same system as the web browser.

To configure the provisioning servers, follow these steps:

1. Edit the `DCM_objects.xml` file.
2. Import the `DCM_object.xml` file into Tivoli Service Automation Manager.
3. Verify that the provisioning servers are created correctly.
4. Verify that the device drivers are created correctly.

5. Verifying that the image discovery configurations are created correctly.

Editing the DCM_objects.xml file

To edit the DCM_objects.xml file, complete the following steps:

1. Open the DCM_objects.xml file in a text editor.
2. Configure the IBM SmartCloud Enterprise account information by editing the icc1.ibm.com server entry as follows:
 - a. Specify the endpoint name to connect to the IBM SmartCloud Enterprise cloud. The default value is IBMCCRAL.

Note: This is the name you fill in Figure 5-5 on page 100.

- ```
<property component="KANAHA" name="name" value="IBMCCRAL" />
```
- b. Specify the name of a security key that is defined in the IBM SmartCloud Enterprise.

```
<property component="KANAHA" name="keyname" value="devkey" />
```
  - c. Specify the user name and password to access the IBM SmartCloud Enterprise account.

```
<sap name="csb" protocol-type="unknown" app- protocol="UNKNOWN" context="NOCONTEXT" port="0" auth-compulsory="false" role="host">
```

```
<credentials search-key="csb" is-default="false">
```

```
<password-credentials username="yourNameHere"
```

```
password="yourPasswordHere" is-encrypted="false" />
```

```
</credentials>
```

```
</sap>
```
3. Configure the WebSphere Cast Iron Appliance information by editing the hci.appliance server entry as follows:
    - a. Specify the WebSphere Cast Iron Appliance host name.

```
<property component="KANAHA" name="hostname" value="" />
```
    - b. Specify the port number to connect to the WebSphere Cast Iron Appliance.

```
<property component="KANAHA" name="port" value="8088" />
```
  4. Save and close the DCM\_objects.xml file.

## **Importing the DCM\_object.xml file into Tivoli Service Automation Manager**

To import the DCM\_object.xml file into Tivoli Service Automation Manager, complete the following steps:

1. Log on to Maximo Start Center.
2. Click **Go To > Service Automation > Cloud Pool Administration**.
3. Click **Import DCM Objects**.
4. In the Select Input File window, click **Browse** to find the DCM\_object.xml file that you modified in the earlier steps, and then click **Open**.
5. Click **OK** to import the Data Center Model objects. A message indicates that the import completed successfully (Figure 5-35).

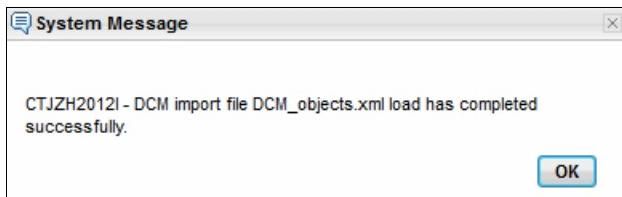


Figure 5-35 Import successful

## **Verify that the provisioning servers are created correctly**

To verify that the import completed successfully and that the provisioning servers were created correctly:

1. Log on to Maximo Start Center.
2. Click **Go To > IT Infrastructure > Provisioning Inventory > Provisioning Computers**.
3. In the Computer field in the Filter area, specify **icc1.ibm.com** and press Enter.
4. Select the **icc1.ibm.com** entry.
5. Click the **Variables** tab, and verify that the information is correct.
6. Repeat steps 2 - 5 to verify the **hci.appliance** entries.
7. Click **Go To > Administration > Provisioning > Resource Pools**.

- Verify that the IBM Compute Cloud Resource Pool is listed as shown in Figure 5-36.

The screenshot shows the 'Provisioning Computers' interface with the title bar 'Provisioning Computers'. The top menu includes 'Bulletins: (0)', 'Go To', 'Reports', 'Start Center', 'Profile', 'Sign Out', and 'Help'. Below the menu is a toolbar with various icons. A navigation bar at the top has tabs for 'List', 'Computer', 'Hardware', 'Software', 'Compliance', 'Recommendations', 'Credentials', 'Workflows', and 'Variables'. The 'Computer' tab is selected. The main area displays a table titled 'Variables' with 7 rows. The table columns are 'Variable', 'Component', and 'Value'. The data is as follows:

| Variable      | Component     | Value                         |
|---------------|---------------|-------------------------------|
| appliance     | Entire system | hci.appliance                 |
| name          | Entire system | IBMCCRAL                      |
| type          | Entire system | ICC                           |
| keyname       | Entire system | devkey                        |
| vlanid        | Entire system | 0                             |
| bootServer    | Entire system | IBM Compute Cloud Boot Server |
| resourcePoold | Entire system |                               |

Figure 5-36 ICC1 computer

Figure 5-37 shows the HCI computer.

The screenshot shows the 'Provisioning Computers' interface with the title bar 'Provisioning Computers'. The top menu includes 'Bulletins: (0)', 'Go To', 'Reports', 'Start Center', 'Profile', 'Sign Out', and 'Help'. Below the menu is a toolbar with various icons. A navigation bar at the top has tabs for 'List', 'Computer', 'Hardware', 'Software', 'Compliance', 'Recommendations', 'Credentials', 'Workflows', 'Variables', and 'Deployments'. The 'Variables' tab is selected. The main area displays a table titled 'Variables' with 3 rows. The table columns are 'Variable', 'Component', and 'Value'. The data is as follows:

| Variable | Component     | Value       |
|----------|---------------|-------------|
| protocol | Entire system | http        |
| hostname | Entire system | 10.7.28.200 |
| port     | Entire system | 8088        |

Figure 5-37 HCI computer

### ***Verifying that the device drivers are created correctly***

To verify that the device drivers were created correctly, follow these steps:

1. Log on to Maximo Start Center.
2. Click **Go To > Administration > Provisioning > Device Drivers**.
3. Verify that the IBM Compute Cloud Discovery Device Driver is listed, as shown in Figure 5-38.

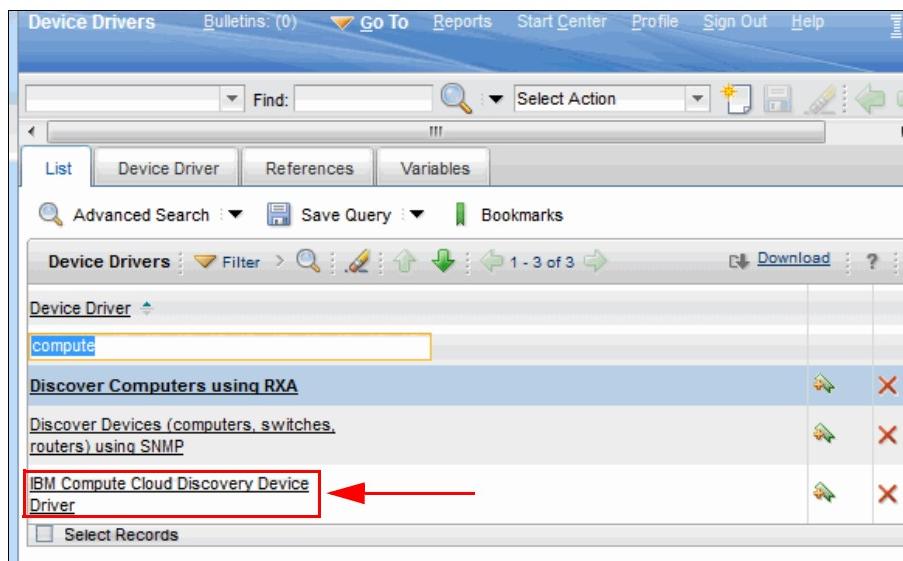


Figure 5-38 Device driver

### ***Verifying that the image discovery configurations are created correctly***

To verify that the image discovery configurations were created correctly:

1. Log on to Maximo Start Center.
2. Click **Go To > Discovery > Provisioning Discovery > Discovery configurations**.

- Verify that the IBM Compute Cloud Image Discovery configuration is listed, as shown in Figure 5-39.

| Discovery Configuration                  | Description                                   | Category           | Method |
|------------------------------------------|-----------------------------------------------|--------------------|--------|
| ibm.com                                  |                                               |                    |        |
| <u>IBM Compute Cloud Image Discovery</u> | Discover Images from IBM Compute Cloud. Other | IBM Compute Driver |        |

Figure 5-39 Discovery configuration

## Creating and enabling cloud pools

The cloud pools are defined in the `vrpool.properties` file.

**Tip:** The `vrpool.properties` file must be located on the same system as the web browser.

A sample `vrpool.properties` file is provided in the `/tmp/hci/config` directory when you expand the `hciExtension.tar` file.

**Upgrade systems information:** You do not need to import the `vrpool.properties` file for an upgrade system. Proceed to step 6 to set the Resource Pool fields.

To import the `vrpool.properties` file into Tivoli Service Automation Manager:

- Log on to Maximo Start Center.
- Click **Go To > Service Automation > Cloud Pool Administration**.
- Click the **Filter Table** icon, next to **Filter**, to display all the resource pools available. Make sure you do not specify any filter. If no resource pools are

displayed, proceed to step 4. If there are resource pools displayed, complete the following steps:

- a. Record the number of resource pools available and also record the **Order** number of the last resource pool ( $n$ ). In this example, it is number 5.

The screenshot shows the 'Cloud Pool Administration' interface. At the top, there are tabs for 'Cloud Pool Overview', 'Cloud Pool Details', and 'Cloud Pool Network Settings'. Below the tabs is a toolbar with various icons for search, filter, and action selection. The main area is titled 'List of Cloud Pools' and displays a table with columns: Order, Cloud Pool Name, Hypervisor Type, and Computer. The table contains six rows, each representing a different resource pool. An arrow points to the fifth row, which is highlighted with a red box and labeled '5 VMC System-p Pool'. The 'Order' column for this row is also highlighted with a red box. The other rows are numbered 0 through 4 and have their respective pool names listed under 'Cloud Pool Name'.

| Order | Cloud Pool Name                   | Hypervisor Type | Computer      |
|-------|-----------------------------------|-----------------|---------------|
| 0     | <a href="#">VMware System x</a>   | VMware          | virtualcenter |
| 1     | <a href="#">System p LPAR</a>     | LPAR            |               |
| 2     | <a href="#">System z pool</a>     | zVM             |               |
| 3     | <a href="#">KVM Local Disk</a>    | KVM             |               |
| 4     | <a href="#">Xen Local Disk</a>    | Xen             |               |
| 5     | <a href="#">VMC System-p Pool</a> | PowerHMC        |               |

Figure 5-40 Existing cloud pool

- b. Edit the `vrpool.properties` file. Change the 0 that precedes all Amazon EC2 Resource Pool entries to one more ( $n+1$ ) than the order number of the last resource pool. Change the 1 that precedes all IBM Compute Cloud Resource pool entries to two more ( $n+2$ ) than the order number of the last resource pool.
- c. Save the `vrpool.properties` file.

4. Click **Import Virtual Resource Pools**. In the Select Input File window, click **Browse** to find the `vrpool.properties` file, and then click **OK**. A message indicates that the import completed successfully (Figure 5-41).

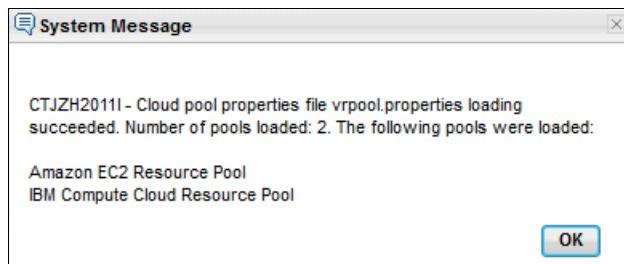


Figure 5-41 Import virtual resource pools

5. Verify that IBM Compute Cloud Resource Pool is listed, as shown in Figure 5-42.

The screenshot shows the "Cloud Pool Administration" interface. At the top, there are tabs for "Cloud Pool Overview", "Cloud Pool Details", and "Cloud Pool Network Settings". Below the tabs is a toolbar with various icons for search, filter, and action selection. The main area is titled "List of Cloud Pools" and displays a table of cloud pools. The table has columns for Order, Cloud Pool Name, Hypervisor Type, and Computer. The data is as follows:

| Order | Cloud Pool Name                 | Hypervisor Type | Computer       |
|-------|---------------------------------|-----------------|----------------|
| 0     | VMware System x                 | VMware          | virtualcenter  |
| 1     | System p LPAR                   | LPAR            |                |
| 2     | System z pool                   | zVM             |                |
| 3     | KVM Local Disk                  | KVM             |                |
| 4     | Xen Local Disk                  | Xen             |                |
| 5     | VMC System-p Pool               | PowerHMC        |                |
| 6     | Amazon EC2 Resource Pool        | KVM             | ec2.amazon.com |
| 7     | IBM Compute Cloud Resource Pool | KVM             | icc1.ibm.com   |

At the bottom of the table, there is a checkbox labeled "Select Records" and two buttons: "Import DCM Objects..." and "Import V".

Figure 5-42 New resource pool

6. Select **IBM Compute Cloud Resource Pool** and set the following fields:

|                       |                                 |
|-----------------------|---------------------------------|
| Resource Pool         | IBM Compute Cloud Resource Pool |
| Target Computer       | icc1.ibm.com                    |
| Image Directory       | /                               |
| Image Repository Host | icc1.ibm.com                    |
7. Then, click **Create New Saved Image Repository** and specify the following fields:

|                      |                   |
|----------------------|-------------------|
| Repository Name      | IBMCoudRepository |
| Target System        | icc1.ibm.com      |
| Saved Image Location | /                 |
8. Set the Saved Image Repository and Instance Image Repository field to the newly created repository name (for example: IBMCoudRepository). Then, click **Save** to save the cloud pool details.
9. Finally, click **Validate** and **Enable Cloud Pool**.

## Creating and registering images

To provision servers in a public cloud, you define a software stack to represent a “real” image in the cloud. Each software stack must be associated with a master image that is then registered to a cloud pool.

When you create a project, or add a server to an existing project, you select a registered image that references an image that is defined in the public cloud.

The Image Discovery automatically discovers images that are available in the public cloud and creates corresponding software stacks.

**Provide a filter condition:** Because the discovery can potentially return hundreds of images, it is mandatory to provide a filter condition to narrow the search to the images that you need.

The example in this section refers to the Red Hat Enterprise Linux 6 (32-bit)(RTP) with the ID 20025200 and the Red Hat Enterprise Linux 6 (64-bit)(RTP) with the ID 20025206.

### ***Defining a filter for IBM Compute Cloud Image Discovery***

To define a filter, you need to edit the Discovery configuration for the IBM Compute Cloud:

1. Click **Go To > Discovery > Provisioning Discovery > Discovery configurations**.

2. Enter IBM Compute Cloud Image Discovery in the Discovery configurations field. Click **IBM Compute Cloud Image Discovery** to open and edit the parameters.
3. Enter the filter that you want to apply to the images.

**Filter string:** A filter is a string. Criteria are defined between square brackets ( [ ] ) and separated by the ampersand (&) character.

The following attributes are supported:

- ▶ IMAGE\_ID
- ▶ IMAGE\_NAME
- ▶ DESCRIPTION
- ▶ VISIBILITY (for IBM Compute Cloud, values are “PUBLIC” or “PRIVATE”).
- ▶ OWNER
- ▶ ARCHITECTURE
- ▶ PLATFORM

For example:

- ▶ To get the private images owned by a person named “steve”, you can add the owner to the filter [VISIBILITY=PRIVATE]&[OWNER='steve'].
- ▶ To specify the exact list of images you want the discovery process to discover and add to the Tivoli Service Automation Manager database, set the attribute IMAGE\_ID with a comma-separated list of IDs [IMAGE\_ID=20025200,20025206].

**Images to avoid:** Do not use any Windows operating system images. Those image types are not supported.

4. Specify the maximum number of images that you want the discovery to import into Tivoli Service Automation Manager. The default value is 10.
5. Save the configuration.
6. Click **Run Discovery**. The **Run Discovery** window opens. In the **Selected Targets** panel, click **Select** and choose **Computers**.
7. Select **icc1.ibm.com** and click **OK**.
8. Click **Submit** to start the discovery process

After the discovery is completed successfully, the software stacks that represent the discovery images are available in Tivoli Service Automation Manager, as shown in Figure 5-43.

The screenshot shows the 'Provisioning Task Tracking' interface. At the top, there's a toolbar with various icons and a status bar indicating 'Bulletins: (0)'. Below the toolbar, there are tabs for 'List' and 'Task', with 'Task' being the active tab. A message box says, 'To see status updates on this page, click the Refresh icon on the toolbar.' Under the message, there are four input fields: 'Provisioning Task' (set to 'IBM Compute Cloud Image Discovery 3/14/12 2:21 AM'), 'Start Date' (set to '3/14/12 02:21:43'), 'Type' (set to 'Discovery'), 'End Date' (set to '3/14/12 02:22:13'), 'Created By' (set to 'MAXADMIN'), and 'Status' (set to 'Success'). Below these fields, there's a section titled 'Provisioning Workflows' with a table. The table has columns for 'Provisioning Workflow', 'Workflow Status', and '% Completed'. One row is visible: 'Discovery.OnDevice' with 'Succeeded' status and a progress bar at 100%. A note below the table says, 'The Status information represents the overall task status. The % Completed information represents the progress of the provisioning task that is running on the targets.'

Figure 5-43 Discovery successful

### **Verifying that the software stacks are created correctly**

To verify that the software stacks were created correctly:

1. Click **Go To > IT Infrastructure > Software Catalog > Software Stacks**.

- Verify that the filtered software stacks are listed as shown in Figure 5-44.

| Software Stack                                          | Version | Vendor |
|---------------------------------------------------------|---------|--------|
| AWSPoolStack                                            |         | AMAZON |
| CDS Depot Stack                                         | 7.2.0.0 | IBM    |
| Cloud Lab Base Agent V2.1 BYOL                          | unknown | IBM    |
| ESXPoolStack                                            |         | IBM    |
| ICCPoolStack                                            |         | IBM    |
| Red Hat Enterprise Linux 6 (32-bit)(RTP)                | unknown | IBM    |
| Red Hat Enterprise Linux 6 (64-bit)(RTP)                | unknown | IBM    |
| Tivoli Common Agent Stack                               | 7.2.0.0 | IBM    |
| Tivoli Common Agent Stack Version 7.2.0.0               | 7.2.0.0 | IBM    |
| Tivoli Common Agent and CDS Depot Stack Version 7.2.0.0 | 7.2.0.0 | IBM    |
| VMware Template – sles11_vs01                           | unknown | VMware |

Figure 5-44 Verify software stacks

### ***Creating a master image***

Next, create a master image for the software stack:

- Click **Go To > IT Infrastructure > Image Library > Master Images**.
- Click the **Add Master Image** icon.
- Go to the **Master Image** tab. Then, specify the following information:
  - In the **Master Image** field, specify the real image name, for example Red Hat Enterprise Linux 6 (32-bit)(RTP).
  - In the **Description** field, specify the real image name, for example, Red Hat Enterprise Linux 6 (32-bit)(RTP).
  - In the **Version** field, specify 2.1.
  - In the **Software Stack** field, specify the real image name, for example, Red Hat Enterprise Linux 6 (32-bit)(RTP).

- Click the **Save** icon to save the master image (Figure 5-45).

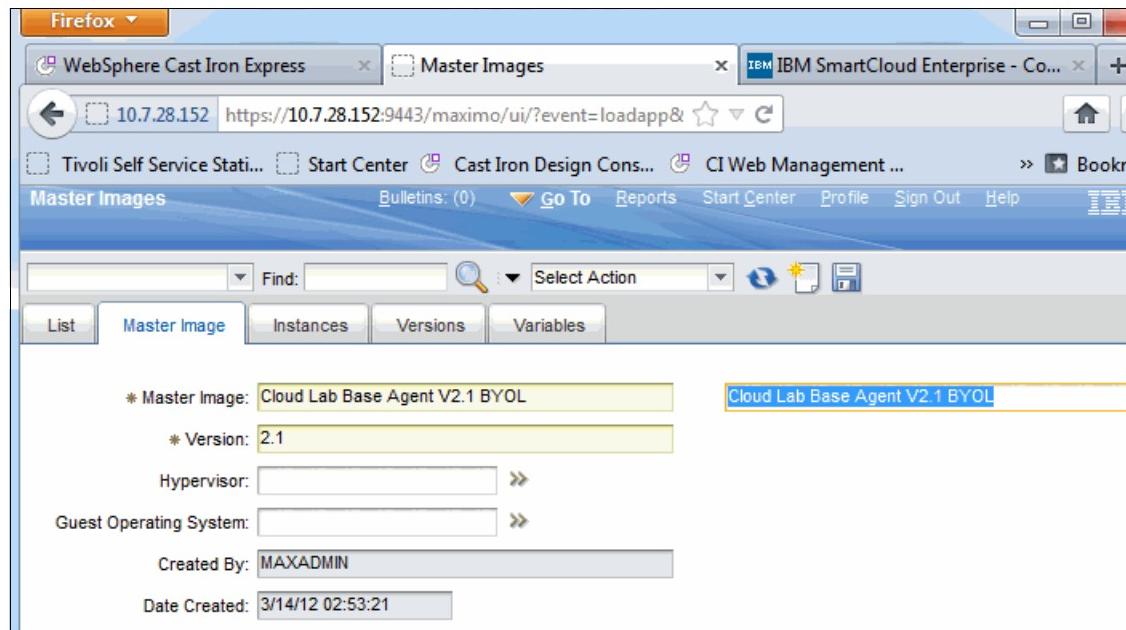


Figure 5-45 Create the master image

### **Registering the master images with the cloud pools**

Register the new master images with the cloud pools to make the images available for selection in the Tivoli Service Automation Manager Self-Service user interface:

- Log on to Tivoli Service Automation Manager as described in 5.3.5, “Logging in to the Tivoli Service Automation Manager” on page 150.
- Click **Request a New Service > Virtual Server Management > Manage Image Library**.
- Register the master images as follows (Figure 5-46 on page 136):
  - Click **Register KVM Image**.
  - In the **Name of Virtual Server Image** field, specify Red Hat Enterprise Linux 6 (32-bit) (RTP) for IBM SmartCloud Enterprise. This value is displayed in the Create Project interface and the Add Server interface. For consistency and ease of deployment, specify the name of the real image to avoid confusion.
  - From the Resource Pool list, select **IBM Compute Cloud Resource Pool**.
  - From the Discovered Image list, select the associated master image.

**Image requirements:** The selected master image must be associated with a real image. The real image must be available in the selected cloud pool, and must be accessible by the account specified by the resource cloud.

Register KVM Image

Register a new KVM server image in the Image Library

**General**

\* Name of Virtual Server Image  
Cloud Lab Base Agent V2.1 BYOL

Description of Virtual Server Image

\* Resource Pool  
IBM Compute Cloud Res

\* Discovered Image  
Cloud Lab Base Agent V2.1 BYOL

**Resources**

|                             | Minimum | Recommended |
|-----------------------------|---------|-------------|
| * Number of Virtual CPUs    | 1       | 1           |
| * Amount of Physical CPUs   | 0.1     | 0.1         |
| * Amount of Memory (in GBs) | 0.051   | 0.051       |
| * Disk Space Size (in GBs)  | 1       | 1           |

OK Cancel

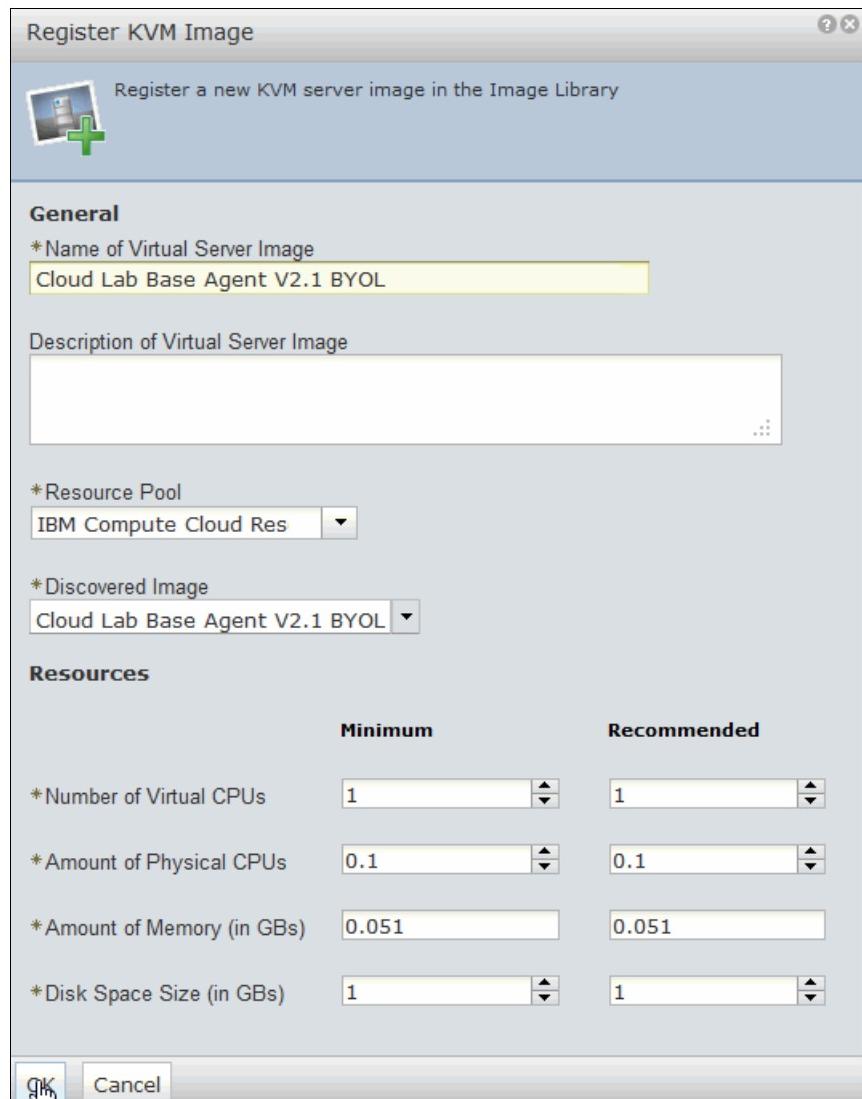


Figure 5-46 Registering the master image

4. Click **OK**. When the registration is complete, the request is marked as Resolved in the My Requests pane, as shown in Figure 5-47.



Figure 5-47 Registering successfully

### ***Creating a team to manage server provisioning***

To create a team to manage server provisioning:

1. Log on to Tivoli Service Automation Manager.
2. Click **Request a New Service -> Virtual Server Management -> Manage Users and Groups**. Then, follow these steps:
  - a. Click **Create Team**.
  - b. In the **Name** field, specify the name of the team.
  - c. From the **Available Users** list, select a user and add the user to the **Selected Users** list, as shown in Figure 5-48 on page 138. Repeat this step until all team members are added.
  - d. Click **OK**.

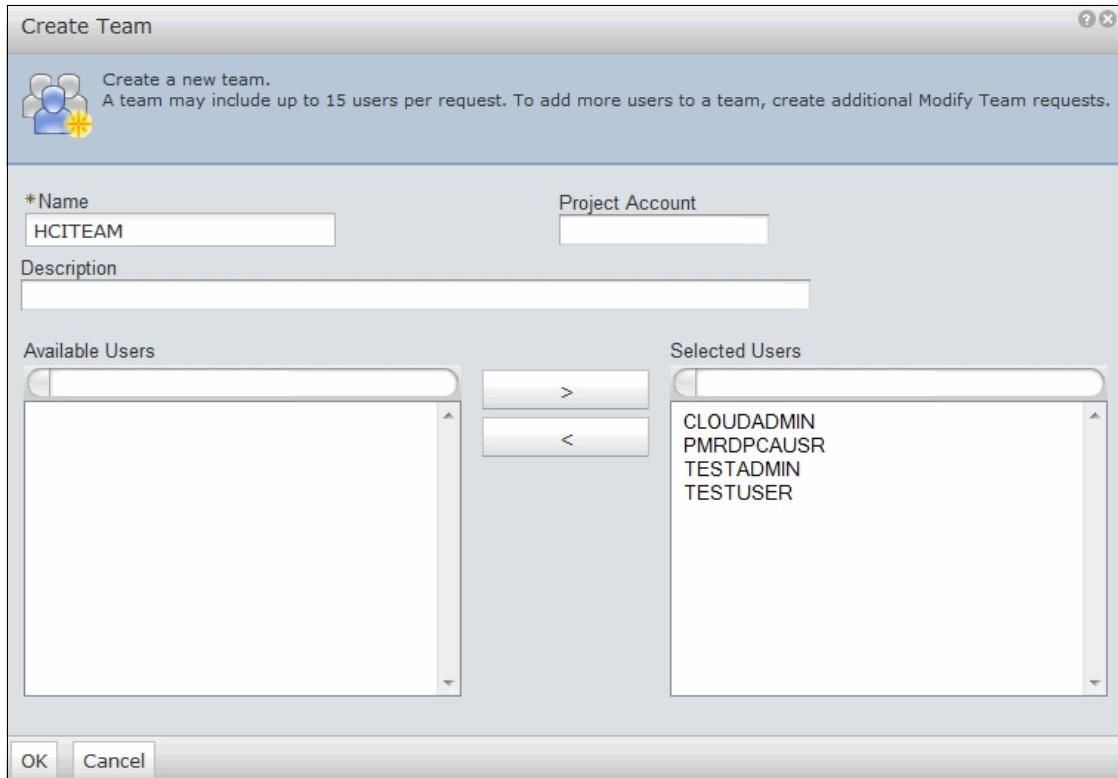


Figure 5-48 Create a Team

When the team is created, the request is marked as Resolved in the My Requests pane.

## 5.3 Scenario usage

The project manager of IBM Redbooks Stores in our scenario had to create new features in applications to support the new stores. The project manager completed the functional tests described in this section to provision servers on the IBM SmartCloud Enterprise.

### 5.3.1 Creating a project with servers in the IBM SmartCloud Enterprise

To create a new project with one server on the SmartCloud Enterprise Cloud and to create the initial project, follow these steps:

1. Log on to Tivoli Service Automation Manager.
2. From the Tivoli Service Automation Manager self - service user interface, click **Request a New Service -> Virtual Server Management**.
3. Click **Create Project with IBM Compute Cloud Servers** as shown in Figure 5-49.

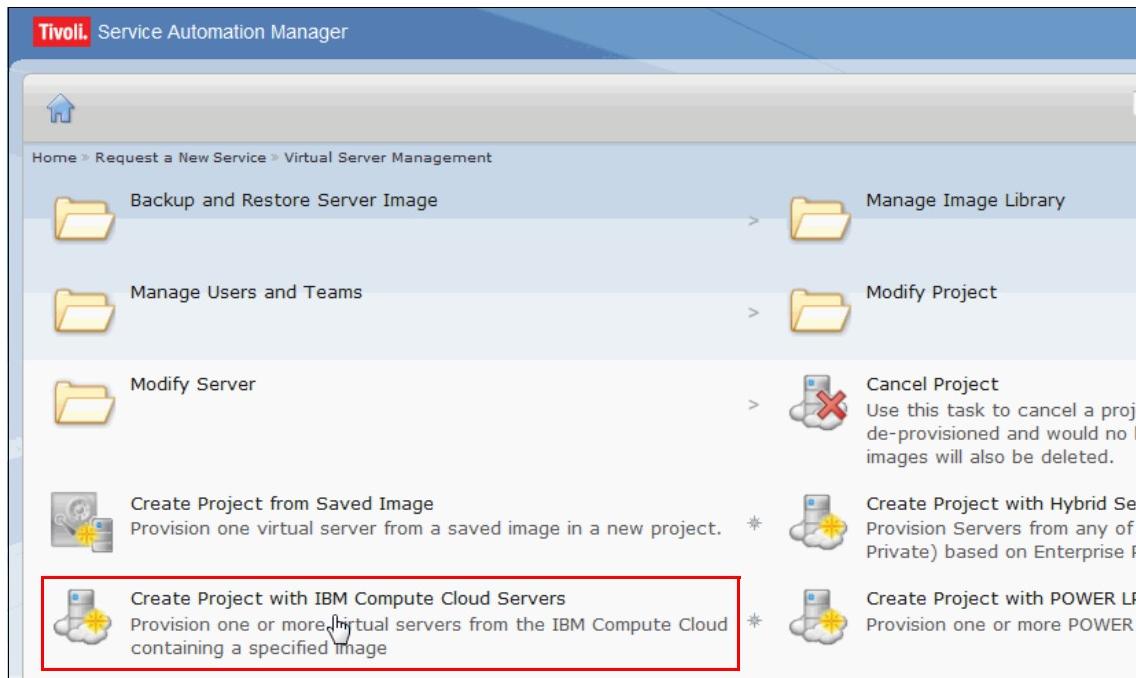


Figure 5-49 SmartCloud Provisioning Service

4. In the Create Project with IBM Compute Cloud Servers dialog box, shown in Figure 5-50 on page 140, complete the following mandatory details:
  - Enter a Project Name and Project Description
  - Select a Team to Grant Access to the project
5. Select **IBM Compute Cloud Resource Pool** from the **Resource Group Used to Reserve Resources** drop-down menu.
6. Select an image to be deployed.

7. Select an instance from the **Select an Instance Type** drop-down menu.
8. Choose the number of servers to be provisioned.
9. Click **OK** to submit the request.

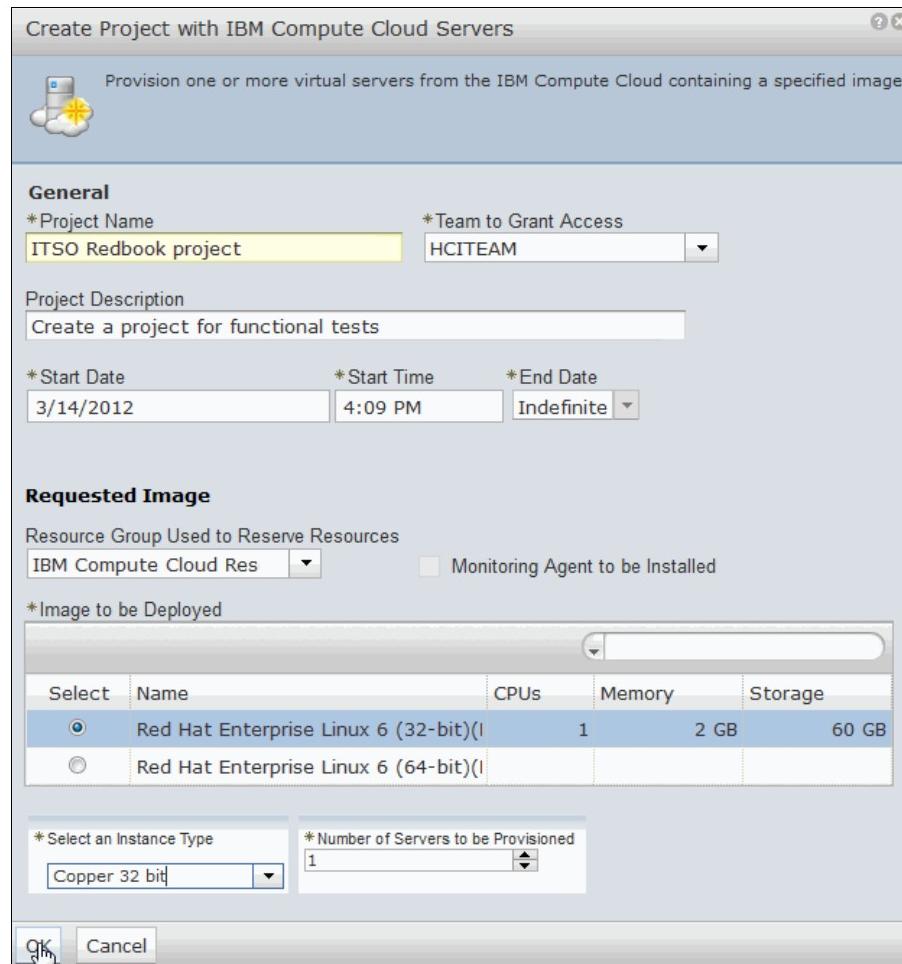


Figure 5-50 SmartCloud provisioning details

The project manager receives an email when the request is in the queue, as shown in Figure 5-51.

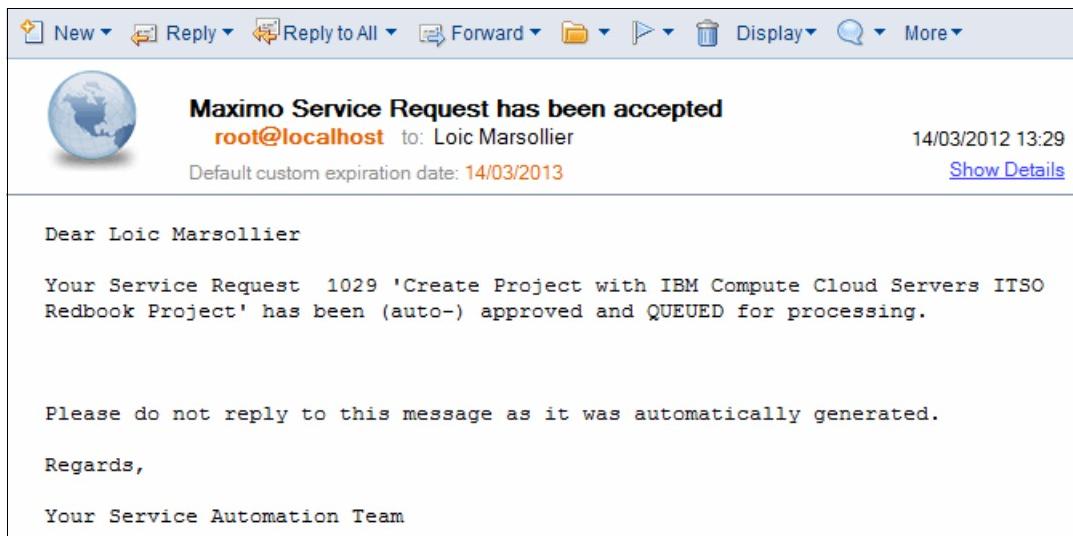


Figure 5-51 Request approved and in the queue

The status of the request is displayed on the My Requests panel on the right side of the user interface. The project is created successfully. One or more servers are provisioned and are running in the IBM Compute Cloud when the status displays as Resolved, as shown in Figure 5-52.

**My Requests**

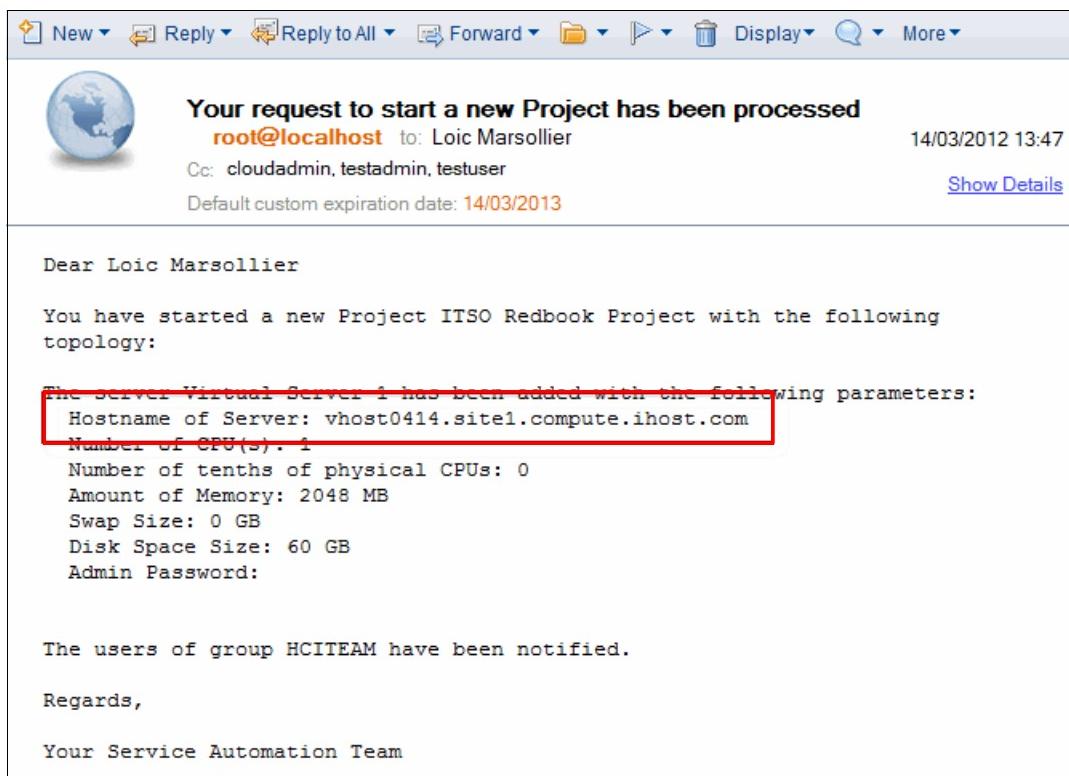
Recent Activity

| Action                                                      | Status   |
|-------------------------------------------------------------|----------|
| Create Project with IBM Compute Cloud Servers HCI Project   | Resolved |
| Register KVM Image Red Hat Enterprise Linux 6 (32-bit)(RTP) | Resolved |
| Register KVM Image Red Hat Enterprise Linux 6 (64-bit)(RTP) | Resolved |
| Create Project with IBM Compute Cloud Servers HCI Project   | Failed   |
| Create Project with IBM Compute Cloud Servers HCI Project   | Failed   |

Show all requests      Manage Requests...

Figure 5-52 My Requests panel

The project manager can view the project status from the My Projects panel on the right side of the user interface. The project manager receives an email with the connection details, as shown in Figure 5-53.



The screenshot shows an email client interface with a toolbar at the top containing 'New', 'Reply', 'Reply to All', 'Forward', 'Display', and other options. The main body of the email contains the following text:

**Your request to start a new Project has been processed**  
root@localhost to: Loic Marsollier 14/03/2012 13:47  
Cc: cloudadmin, testadmin, testuser [Show Details](#)  
Default custom expiration date: 14/03/2013

Dear Loic Marsollier

You have started a new Project ITSO Redbook Project with the following topology:

The server Virtual Server 1 has been added with the following parameters:  
Hostname of Server: vhost0414.sitel.compute.ihost.com

Number of CPU(s): 1  
Number of tenths of physical CPUs: 0  
Amount of Memory: 2048 MB  
Swap Size: 0 GB  
Disk Space Size: 60 GB  
Admin Password:

The users of group HCITEAM have been notified.

Regards,

Your Service Automation Team

Figure 5-53 SmartCloud Enterprise server ready

The first time that the Project Manager goes to the SCE Site, the new server is running and shows as Active (Figure 5-54):

<https://www.ibm.com/cloud/enterprise>

The screenshot shows the IBM SmartCloud Enterprise Control panel. At the top, there's a navigation bar with tabs for Overview, Control panel (which is selected), Account, and Support. Below the navigation bar, there are links for Instances, Images, Storage, and Service Instances, along with a 'View asset catalog' button. The main content area is titled 'My instances' and contains a table of provisioned servers. The table has columns for Instance name, IP Address, Created on, Running, Data center, Total Price/hr, and Status. Two rows are visible: one for 'Andre' (Created on March 7, 2012, Running 8 Days, Data center Boulder1, U.S.A., Status Active) and one for 'Workload VM' (Created on March 14, 2012, Running 1 Hour, Data center Raleigh, U.S.A., Status Active). The 'Workload VM' row is highlighted with a red border. A filter bar at the top of the table allows searching by Tag name, and an 'Add instance' button is located in the top right corner. A note in the top right says 'You are currently using 2 of 4 available instances'.

| Instance name | IP Address      | Created on     | Running | Data center      | Total Price/hr | Status |
|---------------|-----------------|----------------|---------|------------------|----------------|--------|
| Andre         | 170.225.100.145 | March 7, 2012  | 8 Days  | Boulder1, U.S.A. | \$0.125 / UHR  | Active |
| Workload VM   | 170.224.161.159 | March 14, 2012 | 1 Hour  | Raleigh, U.S.A.  | \$0.125 / UHR  | Active |

Figure 5-54 Server provisioned from Tivoli Service Automation Manager

You can click the project to view its content. You can modify the project when the project status is Operational, as shown in Figure 5-55.

**Project: ITSO Redbook Project**

Details for the project and its servers

**Project Details**

|                     |                      |             |
|---------------------|----------------------|-------------|
| Project Name        | ITSO Redbook Project | Operational |
| Project Description | HCI Redbook          |             |
| Project Type        | RDP                  |             |
| Start Date          | 3/14/12 1:29 PM      |             |
| End Date            | Indefinite           |             |
| Team Access         | HCITEAM              |             |
| Total Server(s)     | 1                    |             |
| Active Server(s)    | 1                    |             |

**Servers**

Icon toolbar: Refresh, New, Edit, Delete, Import, Export, Copy, Paste, Find, Filter, Help, Close.

| Select                           | Server Name                       | Management Name | Hypervisor | CPU (%) | Memory (%) | Disk (%) |
|----------------------------------|-----------------------------------|-----------------|------------|---------|------------|----------|
| <input checked="" type="radio"/> | vhost0414.site1.compute.ihost.com | ...             | ICC        | ...     | ...        | ...      |

**Server vhost0414.site1.compute.ihost.com**

|                     |            |               |
|---------------------|------------|---------------|
| <b>General</b>      | <b>CPU</b> | <b>Disk</b>   |
| Hypervisor          | Virtual    | 1 Local 60 GB |
| Provisioning Status | Physical   |               |
|                     | Memory     |               |
|                     | Physical   | 2 GB          |
|                     | Swap       |               |
| Total: 1            |            |               |

**Close**

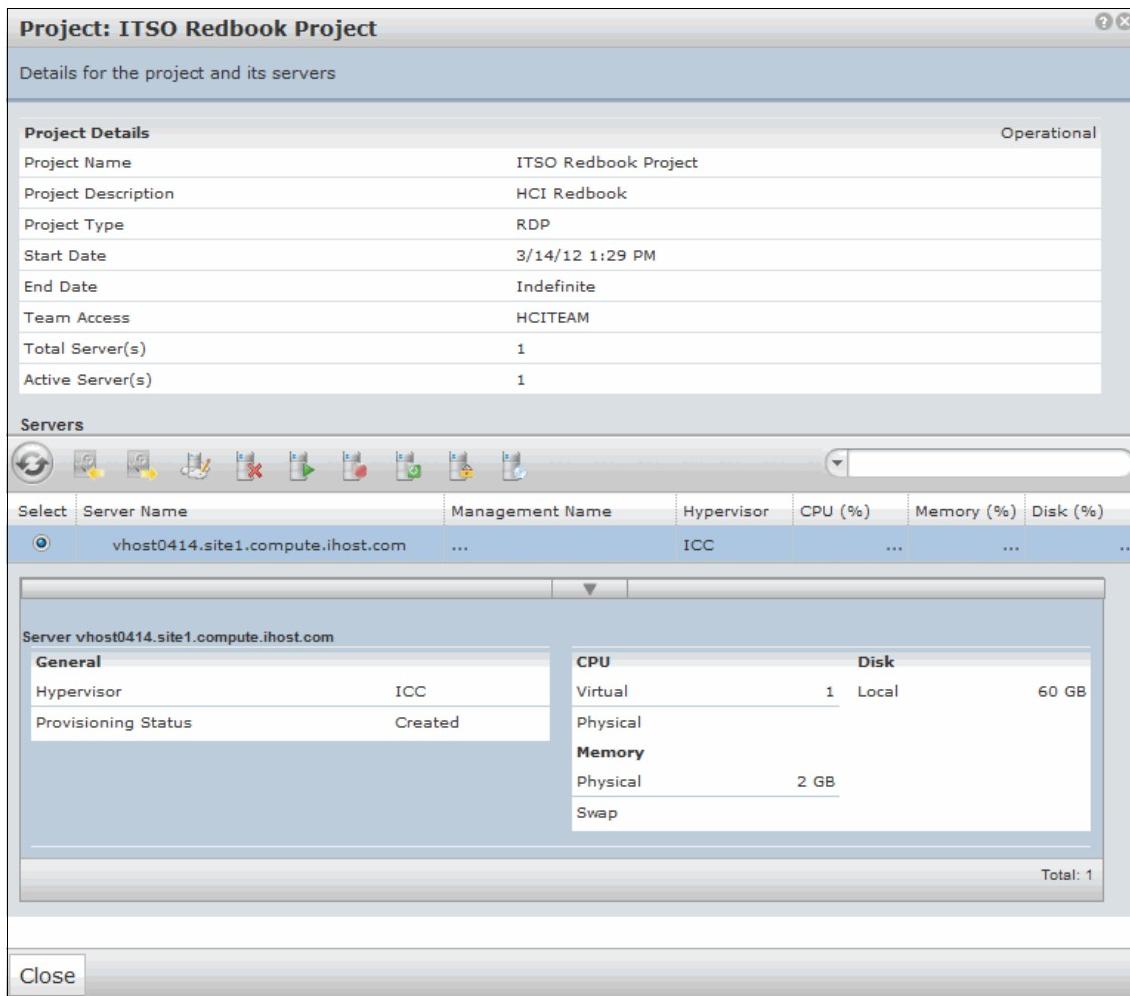


Figure 5-55 ITSO Redbooks Project Details panel

To use this new Linux server, the project manager had to use a client, such as Putty, and the default security certificate that was used when creating the server (Figure 5-56).

The screenshot shows the 'My instances' page with a table of server details. A red box highlights the 'Workload VM' row. Below it, a detailed view for 'Workload VM' is shown with various configuration parameters. An arrow points to the 'Security key pairs' section, which includes a link to 'Instructions'.

| Instance name | IP Address      | Created on     | Running | Data center     |
|---------------|-----------------|----------------|---------|-----------------|
| Andre         | 170.225.100.145 | March 7, 2012  | 8 Days  | Boulder1, U.S.A |
| Workload VM   | 170.224.161.159 | March 14, 2012 | 2 Hours | Raleigh, U.S.A  |

**Workload VM** [ Change name ]

Reboot instance

Instance id : 205002  
OS : Red Hat Enterprise Linux v6  
IP : 170.224.161.159  
Secondary IP(s) : N/A  
Server configuration : Copper - 32 bit (vCPU: 1, RAM: 2 GiB,

Logs :  
→ View logs

Security key pairs :  
→ Instructions

Instance tags  
None  
Add Tag

Figure 5-56 Security instructions

### 5.3.2 Adding one or multiple servers to a project

After a period of time, the first server was not enough to ensure quality tests for the IBM Redbooks Stores. The project manager decided to add a new and larger server to its existing project.

To modify an existing project, follow these steps:

1. From the Tivoli Service Automation Manager self-service user interface, click **Request a New Service -> Virtual Server Management**.
2. Click **Modify Project**.
3. Click **Add IBM Compute Cloud Server** to provision servers on the IBM Compute Cloud.

4. Select a project from the drop-down list. Only projects with the status as Operational are listed.
5. Select a Resource Pool, an Image, and an Instance Type, and the number of servers to be provisioned.
6. Click **OK** to submit the request.

Refer to Figure 5-57.

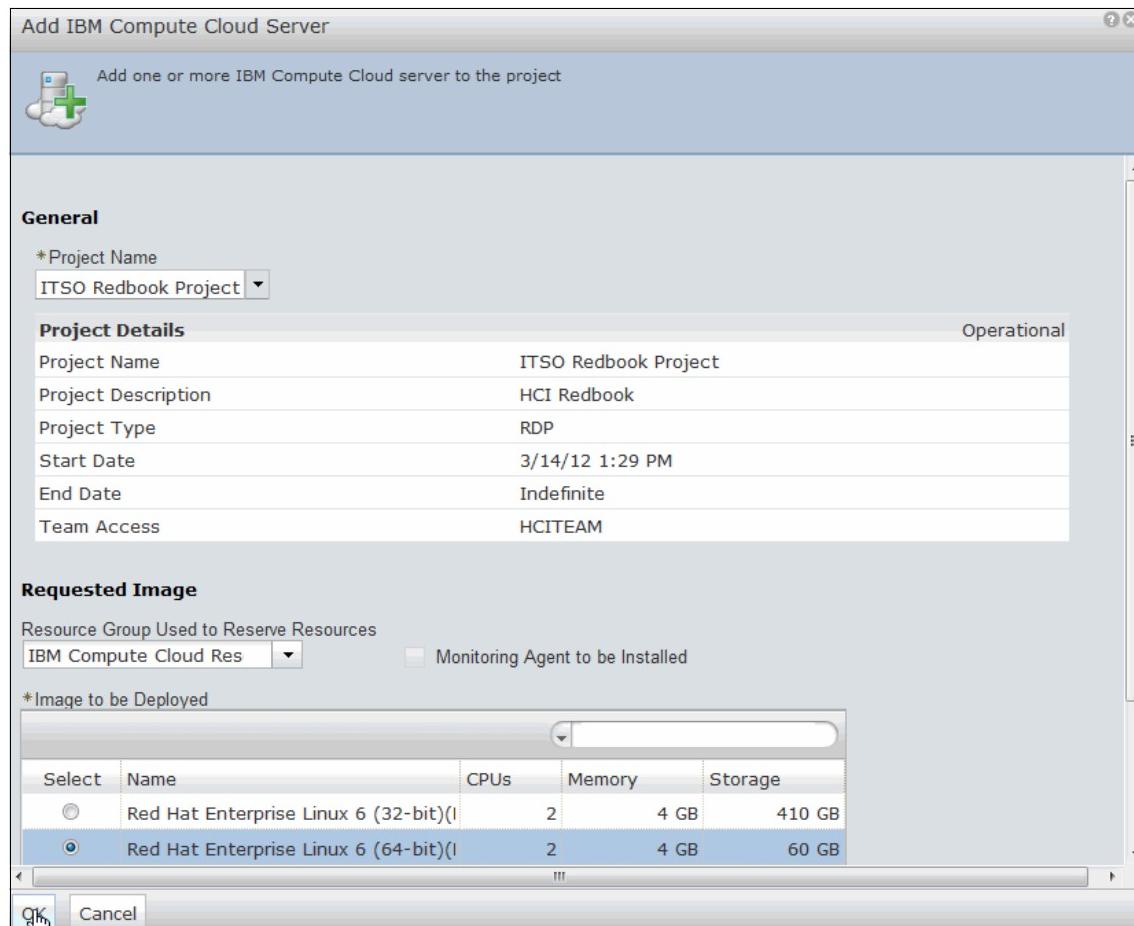


Figure 5-57 Add servers details panel

The status of the request is displayed on the My Request panel on the right side of the user interface. One or more servers are provisioned and running in the IBM Compute Cloud when the status displays Resolved.

### 5.3.3 Removing a server from a project

When the quality tests were finished, the second server was no longer required. To save money, the project manager decided to release this server on the IBM SmartCloud Enterprise.

To remove a server from a project:

1. From the Tivoli Service Automation Manager self-service user interface, click **Request a New Service -> Virtual Server Management -> Modify Project**.
2. Click **Remove Server**. Select the project from the drop-down list. Only projects with the status **Operational** are listed.
3. Select the server from the list of servers of the project (Figure 5-58), and click **OK** to submit the request.

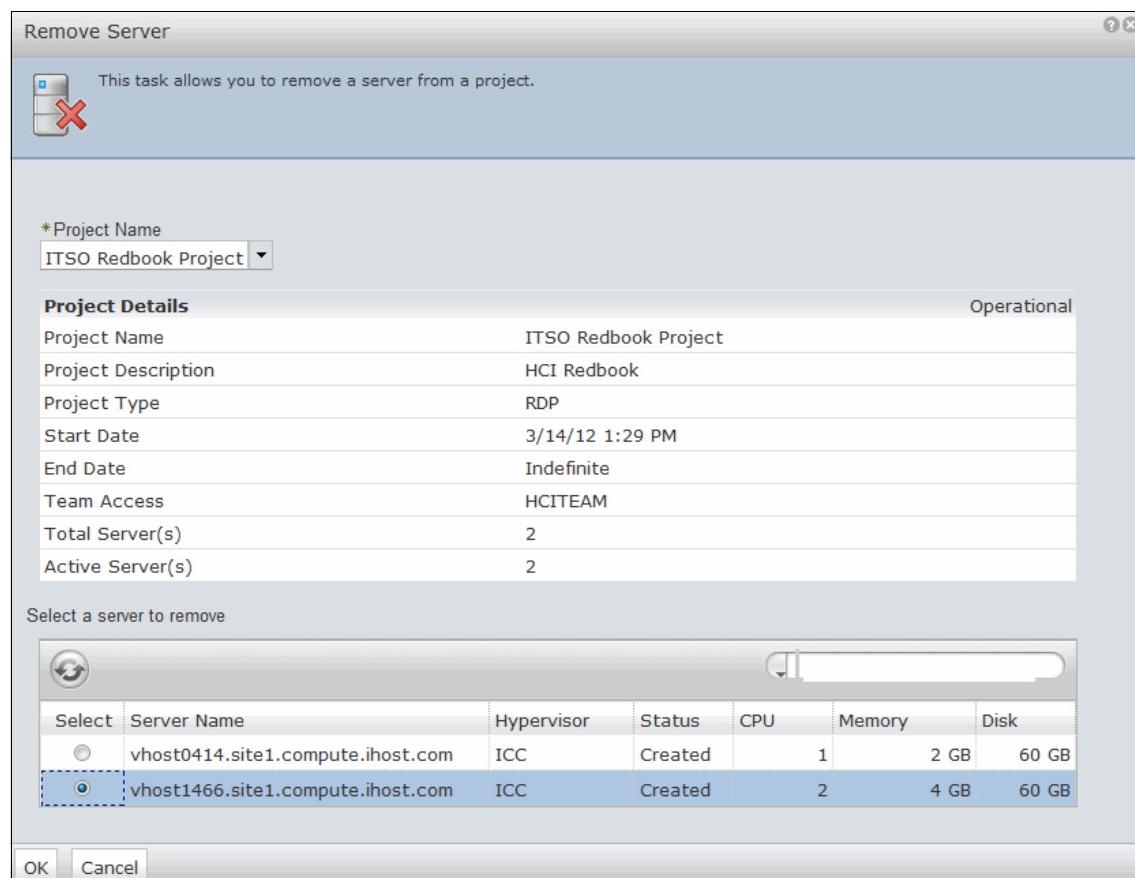


Figure 5-58 Remove Server Details panel

The status of the request displays on the My Request panel on the right side of the user interface. One or more servers is removed from the IBM Compute Cloud cloud when the status displays Resolved.

**Removing the last server:** You cannot remove a server from a project that contains only one server, as shown in Figure 5-59. To remove the last server in a project, cancel the project as described in 5.3.4, “Cancelling a project” on page 149.

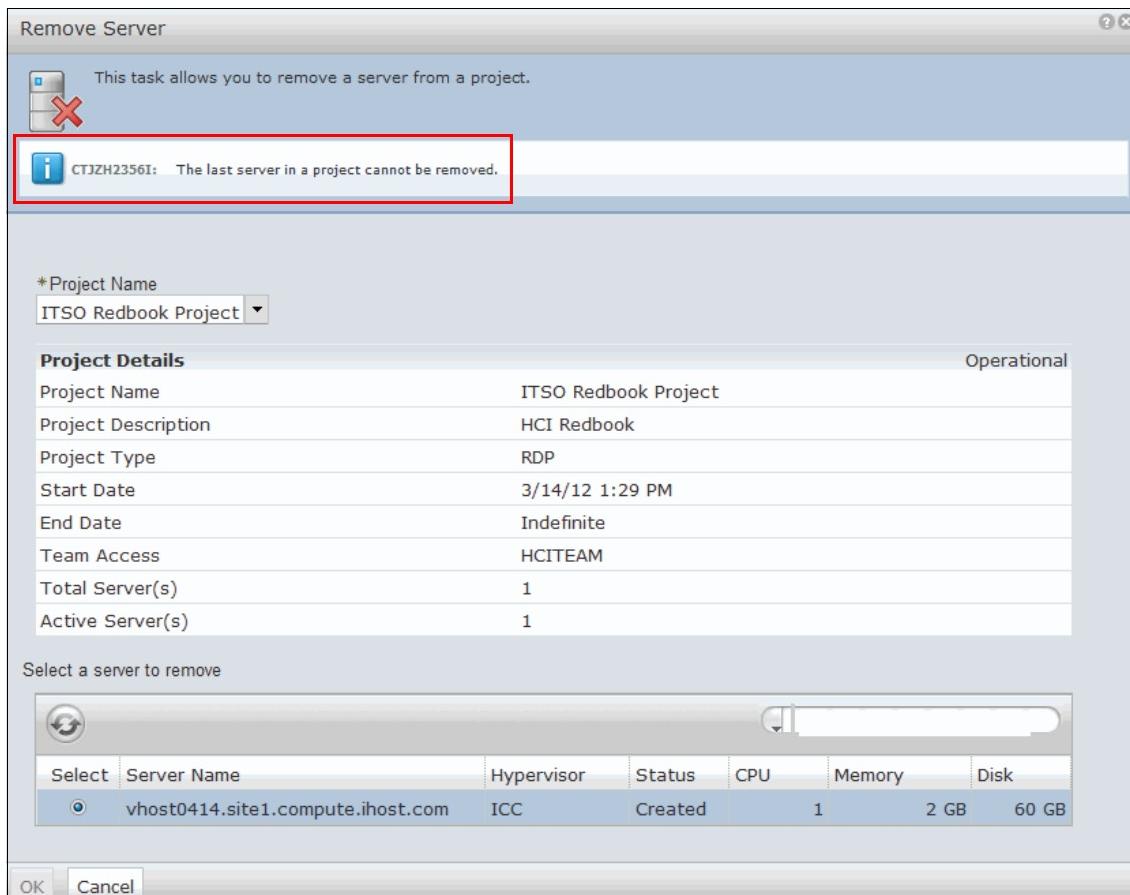


Figure 5-59 Last server

### 5.3.4 Cancelling a project

After the development team qualified the new features of applications, the project manager decided to delete the project environment on the IBM SmartCloud Enterprise.

To cancel a project, follow these steps:

1. From the Tivoli Service Automation Manager self-service user interface, click **Request a New Service -> Virtual Server Management**.
2. Click **Cancel Project**. Select the project to be cancelled from the drop-down list (Figure 5-60).
3. Click **OK** to submit the request.

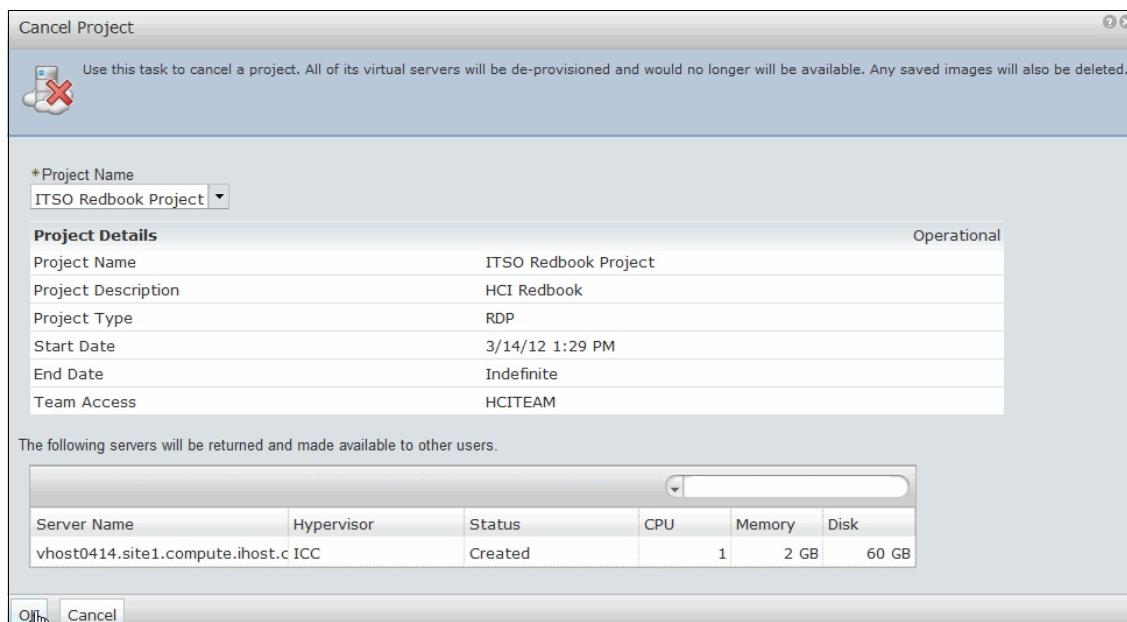


Figure 5-60 Cancel the project

The status of the request is displayed on the My Request panel on the right side of the user interface. All servers are removed from the IBM SmartCloud Enterprise when the status displays as Resolved.

From the My Projects panel on the right side of the user interface, the project status is updated to Decommissioned.

### 5.3.5 Logging in to the Tivoli Service Automation Manager

To login to the Tivoli Service Automation Manager self-service user interface:

1. Open the Tivoli Service Automation Manager self-service user interface:

<https://hostname:9443/SimpleSRM/login.jsp>

2. Log in as the Tivoli Service Automation Manager user. The default user name is PMRDPCAUSR and the default password is maxadmin.

**Browser configuration:** Your browser must be configured to accept cookies. Otherwise, you will not see any action.

### 5.3.6 Logging in to the Maximo Start Center

To login to the Maximo-Administrative user interface:

1. Open the Maximo-Administrative user interface:

<https://hostname:9443/maximo/ui/login>

2. Log in as the Maximo admin user. The default user name is maxadmin and the default password is passw0rd.

## 5.4 Lessons learned

When considering this type of scenario, keep in mind the following points:

- ▶ The Tivoli Service Automation Manager version must be 7.2.1.3 instead of 7.2.1.4 as written in the HCI extension documentation.
- ▶ In the Cast Iron Integration, the source name must be identical in the DCM\_object.xml file, for example, IBMCCRAL.
- ▶ You must have a SmartCloud valid user ID with the ability to create instances.

## 5.5 Summary

The IBM Redbooks Stores company realized the following business benefits and cost savings of the hybrid provisioning implementation:

- ▶ One learning curve
- ▶ One user interface

- ▶ One maintenance point
- ▶ A unique services catalog

IBM Redbooks Stores also realized an improvement in the time to value of its acquisition without increasing its capital expenses by using, on demand, the IBM SmartCloud Enterprise resources and private Cloud Computing Platform.

The development team increased its efficiency with application upgrades and integration in the common applications set. The quick time to market gave IBM Redbooks Stores a competitive benefit over its competitors and made this acquisition more successful.





# Directory Synchronization scenario

In this chapter we continue the IBM Redbooks Store scenario introduced in Chapter 2, “Integrating cloud solutions with Cast Iron” on page 15. Here we detail the directory synchronization portion of that scenario (2.9.3, “Directory Synchronization” on page 37).

In the scenario, the company hires additional people during three months of the year, requiring synchronization of their on-premise user directory with the cloud user directory.

In this chapter we discuss the following topics:

- ▶ Scenario overview
- ▶ Implementing the directory synchronization
- ▶ Running and testing the scenario
- ▶ How the integration works
- ▶ Lessons learned
- ▶ Summary

## 6.1 Scenario overview

In this chapter we show how to synchronize an enterprise Tivoli Directory Server to the public LotusLive. The objective of this scenario is to provide the ability to synchronize user directories from an existing standard enterprise directory service, such as an LDAP server or a compatible directory server like Tivoli Directory Server to a supported public cloud service provider. The result is that on-premise users can access the applications or services both on premise and deployed in the public clouds using on-premise enterprise directories.

Figure 6-1 shows the relationship between all the components across the on-premise and public cloud service provider LotusLive.

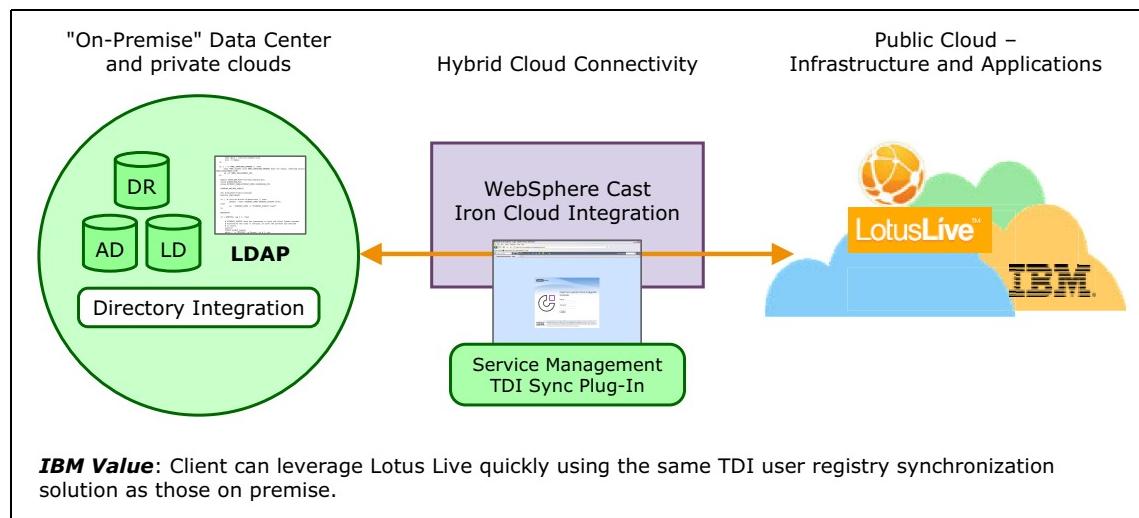


Figure 6-1 IBM Hybrid Cloud integration solution - Directory Synchronization

As previously mentioned, the IBM Redbooks Store has acquired another retail company. After the acquisition the store leveraged a hybrid cloud model to obtain the additional capacity needed for its busy season. During this season, the store will also hire additional staff. Getting the newly hired workers up and running in short period of time requires the provisioning of these additional users in both the on-premise systems and the cloud messaging system.

IBM Redbooks Store uses Lotus Live as the email system for the temporary staff and needs to synchronize its internal LDAP directory, with the temporary staff users, to LotusLive. This synchronization is achieved using the Directory Synchronization bundle installed on the existing Cast Iron appliance and integrating it to the Lotus Live system on the cloud.

In this scenario, the source is an internal LDAP-based server and the target is LotusLive. Both servers are external to the Cast Iron server. Groups and members are defined on the LDAP server.

The general flow of the synchronization process is to take group members from LDAP and add or delete them as subscribers from the LotusLive account. After the members have been added to LotusLive, the new subscribers will receive welcome mail from LotusLive with a URL to set up their new LotusLive account password and other personal details.

Subsequent changes and modifications to member details in the LDAP directory, such as address and phone number, will be synchronized to the LotusLive address book. When a member is removed from an LDAP group, the corresponding user account is also deleted from LotusLive along with any information that is associated with that user.

There are basically two updates going on in the background of LotusLive. One is the update of the address book. The other is the addition and deletion of

accounts under a particular customer of LotusLive. A process flow chart is shown in Figure 6-2.

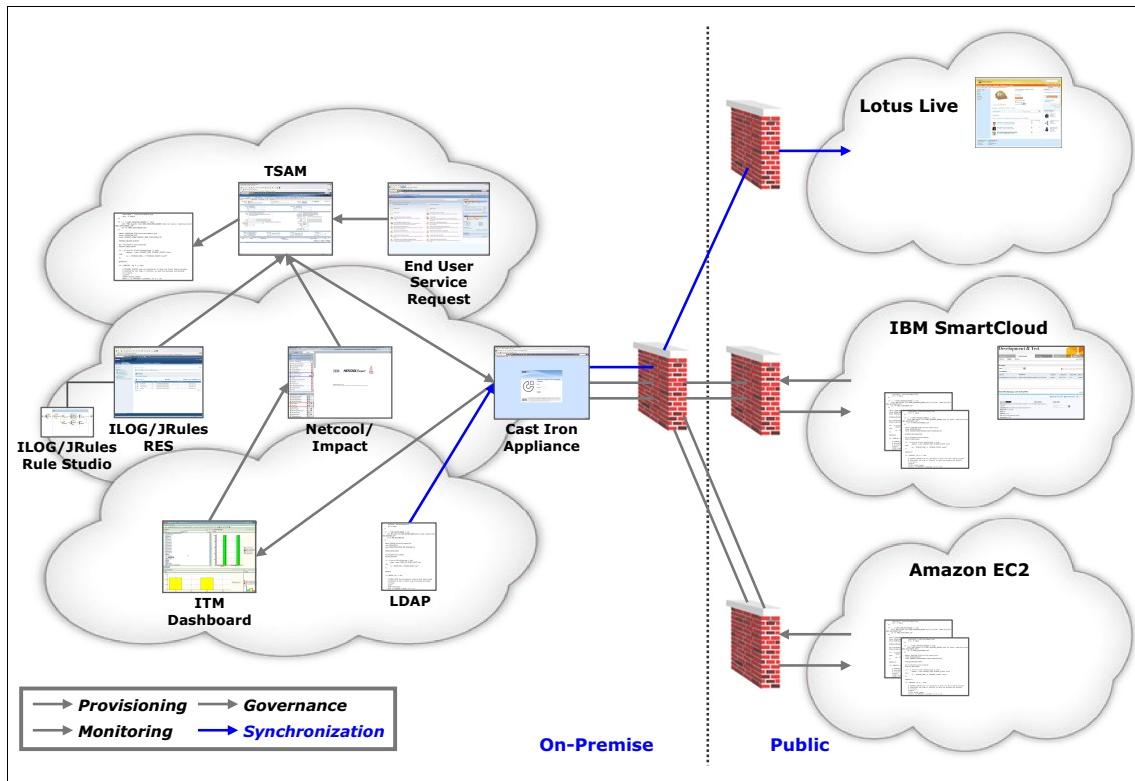


Figure 6-2 Flow chart of the Cast Iron directory synchronization

## 6.2 Implementing the directory synchronization

This section describes how to set up the directory synchronization between the on-premise and cloud directories. It does not cover the full installation of the products. Only the installation steps that are relevant for the scenario are explained.

The IBM products that are used in this scenario include:

- ▶ IBM Service Delivery Manager
- ▶ Tivoli Service Automation Manager
- ▶ IBM Service Management Extensions for Hybrid Cloud 7.2.1

The Service Management Extensions are available at no cost from the IBM Tivoli Software product library. You can download the Service Management Extensions for Hybrid Cloud at the IBM Tivoli Integrated Service Management Library at the following URL:

<https://www-304.ibm.com/software/brandcatalog/ismlibrary/details?catalog.label=ITW1OTS0D>

### 6.2.1 Install Service Management Extensions for Hybrid Cloud plug-ins

In this section, we illustrate and explain in detail the steps needed to install the Service Management Extensions for Hybrid Cloud Plug-ins.

1. Log on to the WebSphere Cast Iron Design Console at the URL:  
<http://hostname/express/>. The default administrator user name is admin and the default password is !n0r1t5@C.

After you are logged in, you will be presented with a window as shown in Figure 6-3 on page 158. In the left navigation pane is the link to **Manage Bundles**. Clicking this link will start the sequence to install new bundles and

upgrade or remove existing bundles. The right pane lists the WebSphere Cast Iron server events that took place recently.

The screenshot shows the Cast Iron Design Console interface. The top navigation bar includes 'Home' and 'Integrations'. The left sidebar has a 'Manage Bundles' section with the sub-instruction: 'Install new bundles, and upgrade or remove existing bundles.' The main content area is titled 'Activity Feed' and displays a list of recent events. The events are listed in descending order of time, starting with 'Integration "BP2LL" updated by sean.' on March 15, 2012, at 4:40:17 PM. Other events include the creation of the same integration on March 15, 2012, at 3:59:43 PM, and the deletion of an integration named 'LDAP Group Members to LotusLive' on the same date at 3:42:18 PM. The activity feed also shows multiple updates to the same integration over the course of the day. At the bottom of the activity feed, there are pagination controls: '1 - 8 of 33' and 'Page 1|2|3|4|5'.

Figure 6-3 Cast Iron Design Console

2. To install the Service Management Extensions for Hybrid Cloud Plug-ins, click **Manage Bundles**.
3. Browse to the folder where the installation package was downloaded, select the Tivoli Directory Integrator Plugin bundle, and click **Upload**.
4. The License Agreement is displayed. To install the plug-in, select **Accept**.

- At the completion of the bundle installation, you will see a list of all bundles that have been installed in the Cast Iron Design Console (Figure 6-4).

The screenshot shows the 'Cast Iron Design Console' interface. At the top, there's a navigation bar with 'Home' and 'Integrations' tabs. Below the navigation is a search bar with a 'Browse...' button and an 'Install New Bundle' button. The main area is a table titled 'Bundle Name' with columns for Version, Vendor, Licence, and Install Date. There are three rows: 'Tivoli Directory Integrator Plugin' (Version 7.2.1.0, Vendor ibm, Licence IBM, Install Date 2012-03-09 20:30:31), 'Provisioning Plugin' (Version 7.2.1.0, Vendor ibm, Licence IBM, Install Date 2012-03-09 20:27:08), and 'Monitoring' (Version 7.2.1.0, Vendor ibm, Licence IBM, Install Date 2012-03-09 20:28:40). The 'Monitoring' row is highlighted with a yellow background.

| Bundle Name                        | Version | Vendor | Licence | Install Date        |
|------------------------------------|---------|--------|---------|---------------------|
| Tivoli Directory Integrator Plugin | 7.2.1.0 | ibm    | IBM     | 2012-03-09 20:30:31 |
| Provisioning Plugin                | 7.2.1.0 | ibm    | IBM     | 2012-03-09 20:27:08 |
| Monitoring                         | 7.2.1.0 | ibm    | IBM     | 2012-03-09 20:28:40 |

Figure 6-4 Tivoli Directory Integrator Plugin has been successfully installed on the Cast Iron server.

## 6.2.2 Create directory entries for WebSphere Cast Iron integration

In this section we detail the steps needed to create user entries on the Directory server for WebSphere Cast Iron integration. Any LDAP standard compatible servers, such as Tivoli Directory Server (TDS), can be used for the integration with WebSphere Cast Iron Directory Synchronization Plug-in. We used the Tivoli Directory Server for this integration scenario.

- Log on to the Tivoli Directory Server Web Administration Tool at the URL: <http://<hostname>:12100/IDSWebApp>. Log in with the directory entry credential that has the authority to create and modify the directory entries, for example, cn=root as we used here, and type in the password (Figure 6-5).

The screenshot shows the 'Tivoli Directory Server Web Administration Tool'. The title bar says 'Tivoli' and 'Tivoli Directory Server Web Administration Tool'. Below it is a 'Directory server login' form. The form has a title 'Enter user name and password'. It contains three fields: 'LDAP Server Name:' with a dropdown menu set to 'TDS', 'User DN:' with a text input field containing 'cn=root', and 'Password:' with a text input field containing '\*\*\*\*\*'. Below the form are two buttons: 'Login' and 'Login to Console admin'.

Figure 6-5 Log in to the Tivoli Directory Server Web Administration Tool

2. Click **Users and Groups** in the left navigation pane (Figure 6-6).

The screenshot shows the Tivoli Directory Server Web Administration Tool interface. The left sidebar has a tree view with nodes like Introduction, User properties, Server administration, etc., and 'Users and groups' is expanded, showing sub-options Add user, Manage users, Add group, and Manage groups. The main content area is titled 'Introduction' and contains a welcome message, a 'Click an item to perform a task.' button, a 'Monitor' section with View Status, View Root DSE, View Logs, and View Cache Status links, a 'Configuration Actions' section with Manage Server Properties, Manage Security, Manage Password Policy, and Manage Replication links, and a 'Content Management' section with Manage Entries link.

Figure 6-6 Tivoli Directory Server Navigation Pane, Users and Groups

3. Click **Add user** and then select the realm from the drop-down menu (Figure 6-7). Then click **Next**.

The screenshot shows the 'Add user' step 1 screen. The left sidebar is identical to Figure 6-6. The main content area is titled 'Add user' and shows 'Step 1' with the sub-step 'AddUser → Step 1 Step 2'. It instructs to 'Select the realm of the new user.' and has a 'Realm:' field with a dropdown menu containing the value 'IBM'.

Figure 6-7 Add user and select a Realm

- Fill in the user information for the entry you are going to create (Figure 6-8).

The screenshot shows the 'Add user' interface in the Tivoli Directory Server Web Administration Tool. The left sidebar navigation includes 'Introduction', 'User properties', 'Server administration', 'Proxy administration', 'Schema management', 'Directory management', 'Replication management', 'Realms and templates', and 'Users and groups' (which is expanded to show 'Add user', 'Manage users', 'Add group', and 'Manage groups'). The main panel title is 'Step 2' under 'AddUser'. It shows 'Step 1' is checked and 'Step 2' is the current step. The 'Realm' field is set to 'IBM'. The 'Naming attribute' section has 'sn' with value 'Loic Marsollier'. Below it, a table lists 'User groups' with 'Marsollier' and 'cn' with 'Loic'. Both fields have an asterisk indicating they are required.

Figure 6-8 Fill in all the required information to create a new user entry in TDS

- After all information required as tagged with an asterisk is filled in, click **Finish** at the bottom of the page.
- You now have the new user entry successfully created in the LDAP directory server. Click **Manage users** in the left navigation pane to see the newly created user entry listed on the right-side information pane (Figure 6-9).

The screenshot shows the 'Manage users' interface in the Tivoli Directory Server Web Administration Tool. The left sidebar navigation is identical to Figure 6-8. The main panel title is 'Manage users'. It shows a 'Select realm:' dropdown set to 'IBM' and a 'View user' button. Below it is a table titled 'Users' with columns for 'Add...', 'Edit...', 'Copy...', 'Delete...', and 'Select Action...'. The table shows one entry: 'Loic Marsollier'. At the bottom, it says 'Page 1 of 1' and 'Total: 1 Filtered: 1 Displayed: 1'. There is also a 'Close' button.

Figure 6-9 A new user entry has been successfully created and listed on TDS directory

7. Repeat the process to create all the user entries you need for integration with the WebSphere Cast Iron Directory Synchronization Plug-in.
8. When you are finished, verify the users have been added by clicking **Manage users** again (Figure 6-10).

The screenshot shows the Tivoli Directory Server Web Administration Tool interface. The left sidebar has a tree view with nodes like Introduction, User properties, Server administration, Proxy administration, Schema management, Directory management, Replication management, Realms and templates, and Users and groups. Under Users and groups, there are links for Add user, Manage users, Add group, and Manage groups. At the bottom of the sidebar is a Logout link. The main content area has a header "Manage users" and a sub-header "Select realm: IBM". It shows a list of users with checkboxes next to their names: Andre Araujo, Chris Rayns, David Kwock, Jag Dhillon, Karen Durston, Loic Marsollier, Sean Liu, and Sunil Joshi. Below the list are buttons for Add..., Edit..., Copy..., Delete..., and other actions like Select, View users, Go, and Filter. At the bottom of the list area, it says "Page 1 of 1" and "Total: 8 Filtered: 8 Displayed: 8". At the very bottom are "Close" and "Logout" buttons.

Figure 6-10 All new user entries are listed and verified in TDS LDAP server

### 6.2.3 Create the WebSphere Cast Iron Directory sync integration

In this section we detail the steps needed to create the integration of WebSphere Cast Iron Director Synchronization with LotusLive public cloud services. The integration relationship needs to be set up and created first on the WebSphere Cast Iron Design Console, before the Directory Synchronization with LotusLive features can become operational.

1. Logon to the WebSphere Cast Iron Design Console at the URL: <http://hostname/express/>. The default administrator user name is admin and the default password is !n0r1t5@C.
2. After you are logged on, click the **Integrations** link at the top of the page. This will take you to the Cast Iron Design Console Integrations page (Figure 6-11)

on page 163). The pane on the right side will display a list of existing integrations and their current status.

The screenshot shows the 'Cast IronDesign Console' interface with the 'Integrations' tab selected. On the left, there's a 'Create' section with dropdown menus for 'Source' and 'Target'. The 'Source' dropdown is currently empty. The 'Target' dropdown has 'LotusLive' listed. On the right, a table lists existing integrations:

| Name                                               | Updated                       | Updated by | Status |
|----------------------------------------------------|-------------------------------|------------|--------|
| ITSO Tivoli Monitoring Server to Off Premise Agent | March 14, 2012 1:46:22 PM GMT | andre      |        |
| TSAM to IBM Compute Cloud                          | March 14, 2012 6:05:13 AM GMT | loic       |        |

Figure 6-11 Cast Iron Design Console Integrations pane

3. In the left pane, under Create, click the Source drop-down menu and select **LDAP Group Members**.
4. In the Target section, select **LotusLive**.
5. After LotusLive is selected, a template for LDAP to LotusLive integration will be available for use. Click the **Use** link next to the template (Figure 6-12).

This screenshot is similar to Figure 6-11 but includes a 'Templates' section at the bottom of the left pane. It shows the 'LDAP to LotusLive' template selected in the dropdown. To its right is a blue 'Use' button with a question mark icon. The rest of the interface is identical to Figure 6-11, showing the 'Create' section and the list of existing integrations on the right.

Figure 6-12 The LDAP Group Members Source and LotusLive Target selected for Integration Creation

6. With the LDAP to LotusLive template selected for integration, the window is split into two halves. The left side is for the source information and the right side is for the target information.
7. We work on the source information entries first:
  - a. Click the blue box on the left side of the window where the number 1 is located.
  - b. Enter the information required to connect to the enterprise LDAP server (Figure 6-13).

Cast IronDesign Console

Welcome back Administrator - IBM

LDAP Group Members to LotusLive

Save Close

**Connection Type**

LDAP Group Members

**LDAP**

**Details**

Name \* TDS

LDAP name \* TDS\_LDAP

LDAP Description Tivoli Directory Server

LDAP URL \* ldap://10.7.28.153:389

LDAP User cn=root

LDAP Password

LDAP Group \* cn=IBM,cn=ibmpolicies,o=ibm.com

Figure 6-13 Fill int all the required information for enterprise LDAP connection

- c. Click **Connect**. If the input is correct you will see a message box at the top of the form confirming that the connect is successfully established with the LDAP server (Figure 6-14 on page 165).

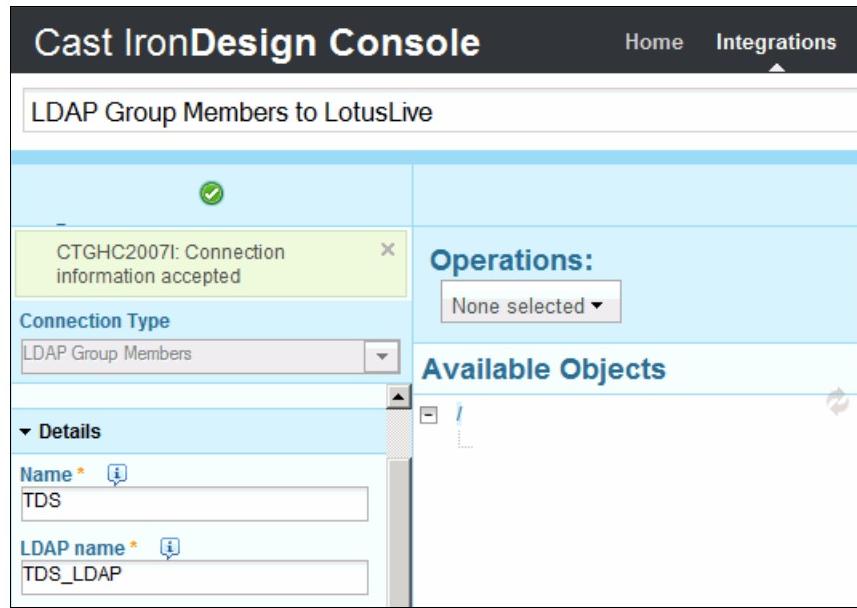


Figure 6-14 Connection to LDAP server has been established successfully

- d. Click **Done** at the bottom right corner of the window to exit from the Source form.
8. Now we move to the Target side.
  - a. Find the orange box at the top the right side of the window where the number 2 is shown, and click anywhere within the orange box.
  - b. A detailed connection information form is presented. Enter the LotusLive account information and the network transfer connection details.

LotusLive uses secure FTP as the network transfer protocol running on port 990.

You will need your LotusLive account name and password for login, along with the customer ID. The LotusLive customer ID is an ID that is assigned to the owner of the account, and which serves as an umbrella of all the users within the group created under this account. It is essential that you enter the correct LotusLive customer ID.

The target information is shown in Figure 6-15.

The screenshot shows the 'Cast IronDesign Console' interface with the title 'LDAP Group Members to LotusLive'. The 'Integrations' tab is selected. The form contains the following fields:

- Connection Type:** LotusLive
- Name \***: LotusLive
- LotusLive name \***: TDS\_LotusLive
- Description**: LotusLive Connection
- FTP Hostname \***: fom.mail.lotuslive.com
- FTP Port \***: 990
- FTP Username \***: ccie@notes.lotuslive.com
- FTP Password**: (Redacted)
- LotusLive Customer ID \***: 20435264

A 'Connect' button is located at the bottom of the form. At the top right, there are 'Save' and 'Close' buttons. The top right corner also displays 'Welcome backAdministrator' and the IBM logo.

Figure 6-15 Provide LotusLive account and connection details in the Target form

- c. Verify that all the required items are correctly entered, then click **Connect**. If everything is correct, a message will be displayed at the top of the form confirming the connection to LotusLive has been established (Figure 6-16 on page 167).

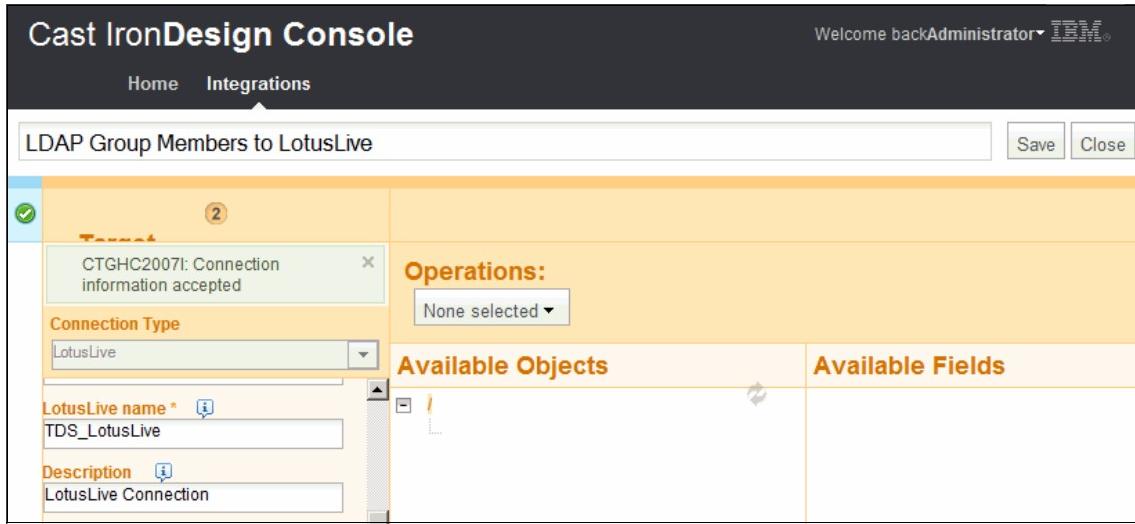


Figure 6-16 The connection to the LotusLive is successfully created

- d. Click **Done** at the bottom right corner of the window to exit from the Target form.
9. Now, with both Source form and the Target form of the integration completed, we need to save them under a single integration. Click **Save** at the top of the Cast Iron Design Console page to save the integration information. You will receive a confirmation box displaying the following statement:

Integration "LDAP Group Members to LotusLive" successfully created.  
Click OK to accept.

Click **OK**.

10. Click **Close** to exit the integration. An overall list of all the integrations that have been defined on the Cast Iron server, including the new Directory Synchronization integration (selected in Figure 6-17), will be displayed.

The screenshot shows the 'Cast Iron Design Console' interface. At the top, there's a header bar with 'Welcome back Administrator' and the IBM logo. Below it, a navigation bar has 'Home' and 'Integrations' tabs, with 'Integrations' being the active one. A sidebar on the left says 'Integrations'. The main area is a table titled 'Integrations' with columns: 'Name', 'Updated', 'Updated by', and 'Status'. There are three rows:

| Name                                               | Updated                       | Updated by | Status  |
|----------------------------------------------------|-------------------------------|------------|---------|
| ITSO Tivoli Monitoring Server to Off Premise Agent | March 14, 2012 1:46:22 PM GMT | andre      | Running |
| LDAP Group Members to LotusLive                    | March 25, 2012 7:30:07 AM GMT | admin      | Stopped |
| TSAM to IBM Compute Cloud                          | March 14, 2012 6:05:13 AM GMT | loic       | Running |

Figure 6-17 List of currently defined integrations with status

## 6.3 Running and testing the scenario

In this section we explain the steps needed to test and run the integration scenario:

1. In the Cast Iron Design Console, go to the integrations window by clicking the **Integrations** link. A list of the current integrations defined on the server and their status (Figure 6-17) will be displayed.

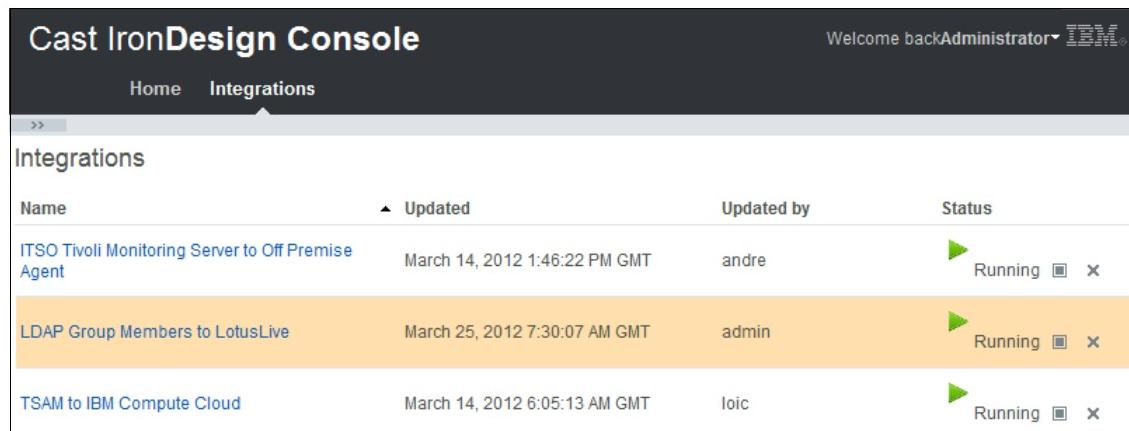
The integration we are going to test and run is called LDAP Group Members to LotusLive. Notice the current status of this integration is Stopped.

2. Activate the integration by clicking the triangle play button next to Stopped.
3. If things are all working properly, a pop-up box displays a confirmation message:

Started integration: LDAP Group Members to LotusLive".

Click **OK** to confirm.

- Verify that on the Integrations window, the integration now displays a status of Running with the triangle play button turning green (Figure 6-18).



The screenshot shows the 'Integrations' section of the Cast IronDesign Console. There are three entries in the table:

| Name                                               | Updated                       | Updated by | Status  |
|----------------------------------------------------|-------------------------------|------------|---------|
| ITSO Tivoli Monitoring Server to Off Premise Agent | March 14, 2012 1:46:22 PM GMT | andre      | Running |
| LDAP Group Members to LotusLive                    | March 25, 2012 7:30:07 AM GMT | admin      | Running |
| TSAM to IBM Compute Cloud                          | March 14, 2012 6:05:13 AM GMT | loic       | Running |

Figure 6-18 The “LDAP Group Members to LotusLive” integration status shows Running

The integration will start to run immediately after the integration has been activated. It will take some time for the Cast Iron Directory Synchronization plug-in to scan the LDAP directory for any addition or deletion updates and transfer the newly created group users to the LotusLive public cloud servers.

A LotusLive process is scheduled to run periodically to process these uploaded directory updates and then populate them into the related LotusLive account.

After the group users are successfully transferred and created within LotusLive account, log on to our LotusLive account for verification (Figure 6-19).

The screenshot shows the 'User Accounts' page in the IBM interface. The left sidebar has sections for Personal (My Account Settings), IBM (User Accounts, Organization Account Settings, Subscriptions, Announcements, Integrated Third-Party Apps, Internal Apps, IBM SmartCloud Labs, Order History, Organization Extensions), and System Settings (Security, Theme, Profile Customization, IBM SmartCloud Directory Options, File Download History). The main area is titled 'User Accounts' and displays a table of users. The table has columns for Name, Role, Status, and Subscription. There are 17 users listed, all pending except for Sean Liu who is active. The subscriptions listed for Sean Liu include SugarCRM™, TeamPoint™, ProjExec® Live - Social Project Management, Salesforce.com™, Skype™, Silanis e-SignDoc™, and Silanis e-SignRoom™.

| Name            | Role               | Status  | Subscription                                                                                                                                    |
|-----------------|--------------------|---------|-------------------------------------------------------------------------------------------------------------------------------------------------|
| Andre Araujo    | User               | Pending | IBM SmartCloud Engage Advanced Trial                                                                                                            |
| Chris Rayns     | User               | Pending | IBM SmartCloud Engage Advanced Trial                                                                                                            |
| David Kwock     | User               | Active  | IBM SmartCloud Engage Advanced Trial                                                                                                            |
| Jag Dhillon     | User               | Pending | IBM SmartCloud Engage Advanced Trial                                                                                                            |
| Karen Durston   | User               | Pending | IBM SmartCloud Engage Advanced Trial                                                                                                            |
| Loic Marsollier | User               | Pending | IBM SmartCloud Engage Advanced Trial                                                                                                            |
| Sean Liu        | Administrator User | Active  | SugarCRM™<br>TeamPoint™<br>ProjExec® Live - Social Project Management<br>Salesforce.com™<br>Skype™<br>Silanis e-SignDoc™<br>Silanis e-SignRoom™ |

Figure 6-19 Users within LDAP server have been transferred and synchronized to the LotusLive account

## 6.4 How the integration works

When the synchronization integration is saved and published, a Tivoli Directory Integration (TDI)-specific configuration file is generated “behind the scenes” that represents the details of that particular integration. It contains the connection parameters and data flow of a user account between LDAP and LotusLive.

After the integration process is activated and starts running, the Directory Synchronization plug-in will constantly check for new updates and changes by comparing the information it receives from LDAP server with a local database it maintains on the server.

Every 15 minutes, it will generate a set of command files that contain all the changes and updates it has detected and upload them to the LotusLive's server by FTP. On the LotusLive server, an integration service will scan its FTP directory every 15 minutes for new command files, and then execute the command files after they are found. When a command file has been processed, a response file is generated by LotusLive and placed in a response directory on the FTP server

At the end of an integration run, the response files are collected and printed to the plug-in's runtime log by the synchronization plug-in. These response files are also downloaded to the Cast Iron server for trace and troubleshooting purposes.

LotusLive provides many services, including email, conferencing, and scheduling. The Directory Synchronization Plug-in does not provision the subscriptions for each individual. These have to be managed by the customer administrator of the LotusLive accounts after the user account have been created or synchronized from the LDAP server. The Directory Synchronization Plug-in will take care of lifetime creation and deletion of LotusLive user accounts based upon an LDAP group.

## 6.5 Lessons learned

In this section, we list common issues you might encounter when implementing the Directory Synchronization Plug-in and the solutions we discovered.

One of the most common problems is not being able to connect to either the LDAP or LotusLive servers. This is usually a connection-oriented problem, which can be the user names and passwords, port numbers, or host IP addresses. A firewall sitting in between the Cast Iron server, the LDAP server, and LotusLive can also create network connectivity issues. If you have a connectivity problem, check the runtime logs. Look for error messages that show unknown host, connection refused or connection timed out. If you cannot resolve the host name of the LDAP server, try using the IP address of the LDAP server. For connection refused messages, make sure you use the correct port number and there is no firewall in between blocking your access.

Another common problem is being unable to find the group name of an LDAP object. The problem can be that the fully distinguished name (DN) of the LDAP group object is typed incorrectly on the integration form, or the URL to access the object is inaccurate. Unknown LDAP object messages or not found messages are typically seen when an LDAP group object you try to read from does not exist in the LDAP directory, or you do not have access to it. You might have provided the wrong user name and password to access an object that is password protected.

## 6.6 Summary

The Service Management Extensions for Hybrid Cloud extends the on-premises suite of products to provide the functionality of unifying user IDs between an enterprise directory and LotusLive.

Customers who want to manage access to off-premise resources and data by using the same enterprise user directory maintained on premise are now able to accomplish and use LotusLive quickly using the same TDI user registry synchronization solution as before. This capability provides better security and control of user access by synchronizing the user directories of on-premise with cloud applications. The automated synchronization means users will be able to gain entry to the information that they have already been authorized to access without human intervention.



# Workload Governance and Management scenario

In this chapter we discuss the Workload Governance and Management scenario that is supported by the IBM Service Management Extensions for Hybrid Cloud. First, we address the functionality of workload governance. Next, we provide the detailed installation and configuration information for the required components for the scenario. Finally, we run the scenario based on a customer use case and provide our results.

In this chapter we discuss the following topics:

- ▶ Scenario overview
- ▶ Installation and configuration
- ▶ Scenario walk-through
- ▶ Lessons learned
- ▶ Summary

## 7.1 Scenario overview

This chapter demonstrates the ability of a hybrid cloud to govern and manage workload. This is done by leveraging several IBM Hybrid Cloud components to create a complete solution that gives a business the ability to apply a business policy to a cloud service and have it enforced. The following IBM components are used to demonstrate this scenario

- ▶ Tivoli Service Automation Manager (TSAM) to administer projects consisting of servers being monitored for performance.
- ▶ WebSphere ILOG JRules for implementing a strategy to provision or decommission resources in the proper location based upon a set of business rules
- ▶ IBM Tivoli Monitoring (ITM) to collect monitoring data (CPU utilization %) of both on-premise and cloud instances
- ▶ Netcool/Impact and Event Correlation Service (ECS) to analyze CPU utilization, detect and correct abnormal situations

Figure 7-1 shows a high level view of this solution.

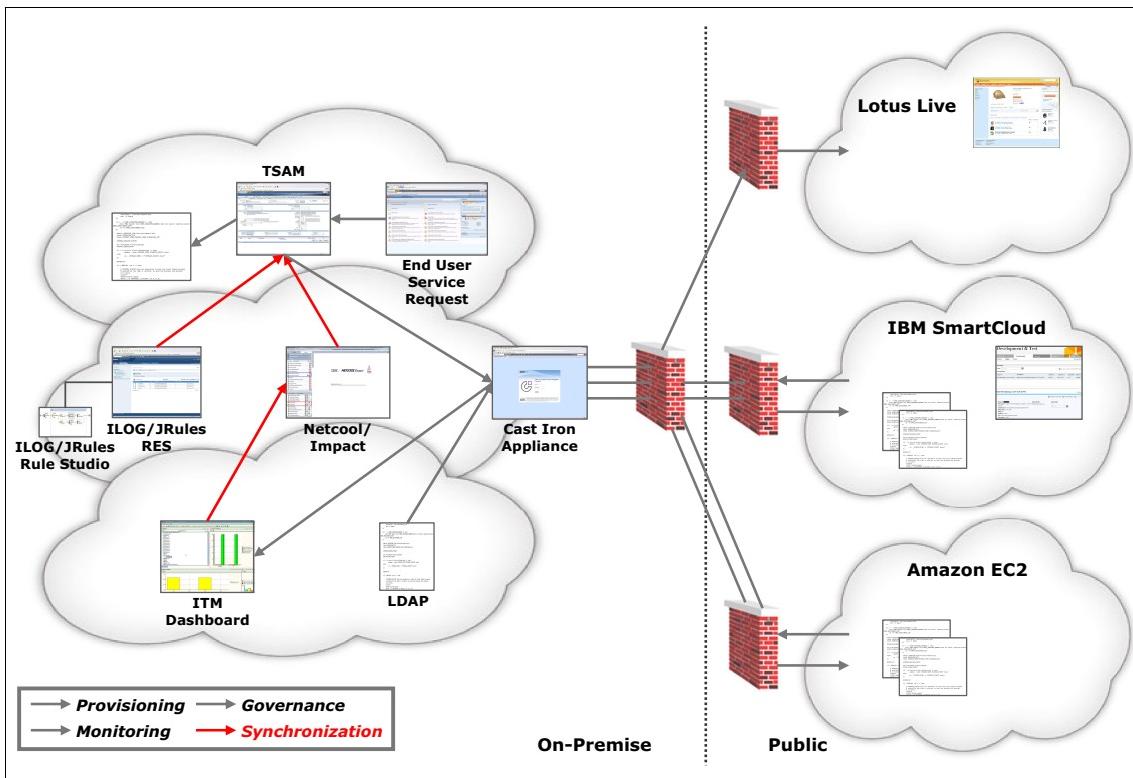


Figure 7-1 Governance Scenario Overview

### 7.1.1 Tivoli Service Automation Manager

The Workload Governance and Management scenario leverages Tivoli Service Automation Manager (TSAM) to offer a new TSAM service catalog offering called “Create Project with Hybrid.” This offering allows for a project that creates resources both with on-premise and public cloud resources based on business rules loaded in the ILOG Rules Execution Server.

### 7.1.2 ILOG JRules Execution Server

After TSAM creates the request, it passes the service request information to a Transparent Decision Service on the ILOG JRules Execution Server, which invokes the business rules to determine where to provision the requested resources. The resource location is returned to TSAM and TSAM continues to process the request in the cloud as indicated by the Decision Service.

### 7.1.3 Tivoli Monitoring

After the resources are provisioned Tivoli Monitoring continues to monitor the resources that are part of the request project and sends data to the Netcool Impact (Event Correlation Service). Impact will determine if resources need to be moved based on business rules.

### 7.1.4 Netcool Impact

After Tivoli Monitoring Data is received about the different components of the hybrid service, the data is processed in Netcool Impact to determine if any additional resources need to be provisioned to stay within business policy. The Event Correlation Service in Netcool determines whether resources need to be allocated or deallocated. If any action is needed for a given cloud service, Netcool Impact puts in a service request to TSAM to perform the action.

## 7.2 Installation and configuration

In this section we explain the prerequisites, installation, and configuration of the workload governance and management scenario. In addition, we cover the environment architecture of our lab scenario.

### 7.2.1 Prerequisites and installation

Table 7-1 shows the software components that are required for the workload governance and management scenario.

*Table 7-1 Products needed for Workload Governance and Management scenario*

| Product                    | Version | Description                                                                                                           |
|----------------------------|---------|-----------------------------------------------------------------------------------------------------------------------|
| Cast Iron Operating System | 6.1.0.2 | Enables companies to rapidly connect their hybrid world of public clouds, private clouds, and on-premise applications |

| Product                           | Version         | Description                                                                                                                                                                  |
|-----------------------------------|-----------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Tivoli Monitoring                 | 6.2.2 FixPack 3 | Manages system and network applications, tracks the availability and performance of your enterprise system, and provides reports to track trends and troubleshoot problems   |
| Tivoli Netcool/Impact             | 5.1.1 FixPack 1 | Streamlines event and alert management, business service management, and incident and problem management processes                                                           |
| Tivoli Service Automation Manager | 7.2.2           | Helps to automate the provisioning, management and de-provisioning of IT resources, hardware server, networks, operating systems, middleware, and application-level software |
| WebSphere ILOG JRules             | 7.1.1           | Provides functionality to build and deploy rule-based applications for Java, mainframe and service oriented architecture-based environments                                  |

To support the solution, the following steps are assumed to have taken place:

1. Install and configure the Tivoli Service Automation Manager software.
2. Deploy the IBM WebSphere Cast Iron Appliance.
3. Install the Service management Extensions for Hybrid Cloud.
4. Contact IBM SmartCloud Enterprise service to obtain a user account.
5. Configure each hybrid service.

To demonstrate this scenario an IBM x3550 and an IBM LS42 Blade, each with 32 GB of RAM, was used. Each server had 1 TB of Disk storage to host the

images required for this scenario. Figure 7-2 is a logical diagram of the configuration.

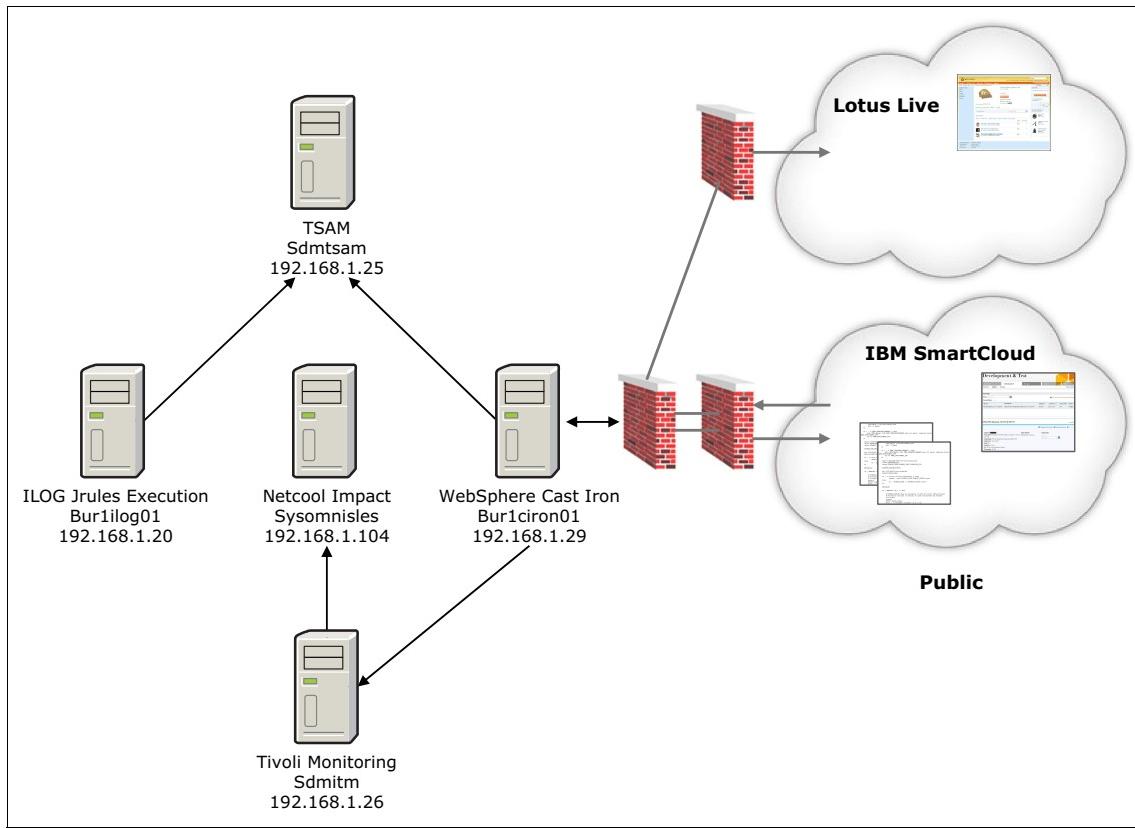


Figure 7-2 Lab environment used to demonstrate Workload Governance scenario

The remaining configuration steps discussed in this section are specific to the Workload Governance and Management scenario. We assume that you have already completed the configuration documented in Chapter 4, “Monitoring scenario” on page 61 and Chapter 5, “Provisioning scenario” on page 93. The configuration documented in this section builds on top of the work that is done in the Monitoring and Provisioning scenarios.

## 7.2.2 TSAM additional configuration

In this step, you will reactivate the PMRDPSIVIEW:

1. Stop the TSAM Server.

2. On the administrative system, log on as root (for UNIX) or as an administrative user (for Windows) and open the command window.
3. Change directory to the Maximo home directory.
4. Run the following script to import the PMRDPSIVIEW configuration:
 

```
./importObject.sh
-input=../en/rdp_pmp/V750_06.content/RDP_pmrdfsview_chg.xml
```
5. Run the following script to import the PMRDPSIVIEW object:
 

```
structure./importObject.sh
-input=../en/rdp_pmp/V750_06.content/NBAPITSAM_maxintobject_change.xml
```
6. Run the following script to commit the changes to the database:
 

```
./configdb.sh
```
7. Restart the Tivoli Service Automation Manager server.

### **7.2.3 Deploying the Event Correlation Service on the Tivoli Netcool/Impact server**

In this section we explain the process used to deploy the event correlation service on the Tivoli Netcool/Impact server. The Event Correlation Service is delivered as a single compressed file called EventCorreelationService.tar, which is part of the download from the IBM Service Management Library.

1. Log on to the Tivoli Netcool/Impact host as the netcool user.
2. Transfer the EventCorrelationService.tar file from the original media to the /tmp directory.
3. Change to the /tmp directory.
4. Extract the contents of the EventCorrelationService.tar file to create an ecs directory that contains all the necessary files.
5. Change to the ecs directory.
6. Define and export the NCHOME Environment variable to specify the location of the Tivoli/Netcool products. The default vale for Tivoli Netcool v5.x is /opt/ibm/netcool1.
7. Run the installation script: `./setupEcs.sh T`

### **7.2.4 Editing the Event Correlation Service configuration file**

The next step is to edit the event correlation service configuration file.

1. Log on to the Tivoli Netcool/Impact host as the netcool user.

2. Change directory to the ecs directory.
3. Edit the ecs.properties file to replace each of the following variables with the correct value for your system.

*Table 7-2 Parameters to modify in ecs.properties*

| Variable   | Description                                                                 | Example          |
|------------|-----------------------------------------------------------------------------|------------------|
| <TSAMHOST> | The host name or IP Address of the Tivoli Service Automation Manager Server | sdmsam.dlk.local |
| <TSAMPORt> | The port used for REST APIs by Tivoli Service Automation Manager            | 9443             |
| <TSAMUSER> | The user ID to connect to Tivoli Service Automation manager                 | maxadmin         |
| <TASAMPWD> | The password for the Tivoli Service Automation Manager user ID              | password         |

### 7.2.5 Installing the WebSphere ILOG JRules Rules Execution Server

Use the following steps to install the ILOG Rules Execution Server:

1. Log on to the proposed Rules Execution Server host as the user who will run the Rules Execution Server.
2. Ensure that the Java Virtual machine is in your execution path. The Rules Execution Server requires an IBM 32-bit JVM.
3. Download the JRules\_V711\_CZLX6ML.bin file from IBM Passport Advantage®. In our case we downloaded it to the /software directory.
4. Change to the /software directory and start the ./JRules\_V711\_CZLX6ML.bin file.
5. In the installation wizard, select the appropriate language and click **OK**.
6. In the Introduction window, click **Next**.
7. Accept the terms of the license agreement and click **Next**.
8. Select the **Rules Execution Server for Java** check box. Click **Next**.
9. Choose the installation location. Click **Next**.
10. Select the JVM to be used by the Rules Execution Server and click **Next**.

11. Review the pre-installation summary and click **Install**.
12. Click **Done** to exit the installer.

### 7.2.6 Verifying the installation of the Rules Execution Server

To validate the Rules Execution Server is installed correctly, follow these steps:

1. Change to the Rules Execution Server Installation directory.
2. Start the Rules Execution Server by launching the sales console.
3. Connect to the Rules Execution Server:  
`http://server:8080/res`
4. Login with the resAdmin user ID.

### 7.2.7 Deploying the hybrid cloud RuleApp

In this section we explain the deployment of the hybrid cloud ruleapp.

1. Copy the `hciRuleapp.jar` file to the `/software` directory.
2. Login to the Rules Execution Server at `http://server:8080/res` with the user `resAdmin`.
3. In the Rules Execution Server console, click the **Explorer** tab.
4. In the RuleApps View pane, click **Deploy RuleApp Archive**.
5. In the Deploy RuleApp Archive window click **Browse**, specify the location of the ruleapp file, select **Increment RuleApp major version**, and click **Deploy**.

### 7.2.8 Additional Tivoli Monitoring configuration

You cannot govern and manage workloads on cloud instances that have only been configured for monitoring. You also cannot migrate projects to support workload governance after projects have been created. Currently, we are using cloud instances only configured for monitoring. The following steps are required to support the workload governance scenario:

1. Log on to the Linux IBM SmartCloud Enterprise instance as the default user `idcuser`.
2. Use the **sudo bash** command to start a shell session as the root user.
3. Change directory to the operating-system startup directory.
4. Modify the `/etc/init.d/ITMAgents1` file that was generated during the installation process to enable the workload governance scenario to be

deployed. You need to enable the root user to start the Tivoli Monitoring agents and change the host name reported by the cloud interface.

The original entry is shown in Example 7-1.

---

*Example 7-1 ITMAgents1 file before update*

---

```
start_all() {/bin/su - root -c "/opt/IBM/ITM/bin/itmcmd agent start
1z >/dev/null 2>&1" }stop_all() {/bin/su - root -c "
/opt/IBM/ITM/bin/itmcmd agent stop 1z >/dev/null 2>&1" }
```

---

The edited entry is shown in Example 7-2.

---

*Example 7-2 Edited ITMAgents1 file*

---

```
start_all() { configFile=$CANDLEHOME/config/lz.ini myName=`hostname
-f` cp $configFile /tmp/lz.ini.bck cat /tmp/lz.ini.bck | sed
"s/^CTIRA_HOSTNAME=.*$/CTIRA_HOSTNAME='$myName'/g" | sed
"s/^CTIRA_SYSTEM_NAME=.*$/CTIRA_SYSTEM_NAME='$myName'/g" >
$configFile if [`id -u` -eq 0]; then $CANDLEHOME/bin/itmcmd agent
start 1z >/var/log/itmcmd.log 2>&1 else sudo $CANDLEHOME/bin/itmcmd
agent start 1z >/var/log/itmcmd.log 2>&1fi }stop_all() { if [`id
-u` -eq 0];
$CANDLEHOME/bin/itmcmd agent stop 1z >/dev/null 2>&1 else sudo
$CANDLEHOME/bin/itmcmd agent stop 1z >/dev/null 2>&1 fi }
```

---

## 7.3 Scenario walk-through

In this scenario, a retail company has an inventory system that leverages two application servers. One application server runs at 80 percent CPU utilization and the other application server runs at 20 percent utilization. This produces a system that averages 50 percent CPU utilization. To handle peaks in workload, the company has implemented a policy that ensures that CPU utilization for the application servers should stay between 42 and 75 percent. This means that during normal workload, the two application servers are adequate because between the two application servers they average 50 percent CPU utilization.

During the holiday season, though, the two application servers start running at 100 percent CPU utilization. This would normally cause customers not to be serviced. However, because this retail company has implemented the policy that CPU utilization should stay between 42 percent and 75 percent, an additional server is automatically provisioned. With the addition of this third server, the CPU utilization drops from 100 percent CPU utilization between the two application servers to 67 percent CPU utilization between the three application servers. This allows all customers to be serviced.

After the holiday season, the first application server drops back down to 80 percent CPU utilization. The second application server drops back down to 20 percent CPU utilization. The third application server drops to 0 percent CPU utilization. This causes the overall CPU utilization to drop to 33, which triggers the automatic de-provisioning of the third application server based on the policy of CPU utilization should stay between 42 and 75 percent. After the third application server is de-provisioned, the CPU utilization drops back to 50 percent between the two application servers and within the policy of CPU utilization should be between 42 and 75 percent.

### 7.3.1 Resources added to cloud project to support business policy

In this section we demonstrate a resource being added to an application based on peak season load. As you will see in this scenario, the two servers that make up the cloud service can no longer handle the workload based on the policy that CPU usage must stay between 42 and 75 percent. To enforce this policy, the workload governance and management scenario triggers the deployment of additional resources (a third compute resource) to ensure the cloud service stays within the business policy.

First we look at the cloud resource project resource details (Figure 7-3) in TSAM to see the details of the cloud resource that has been provisioned to support the retail application.

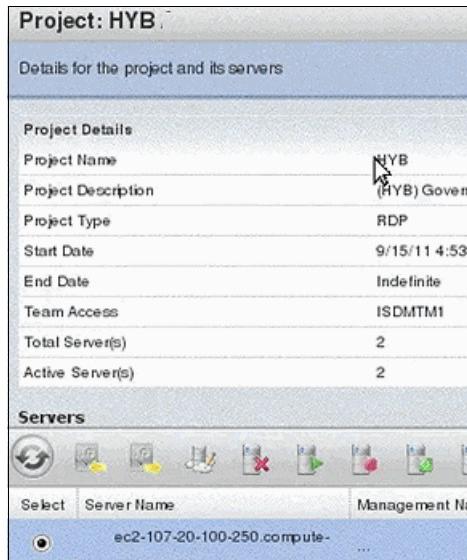


Figure 7-3 Project details

Then, from TSAM, we see the hybrid project has been successfully deployed (Figure 7-4).

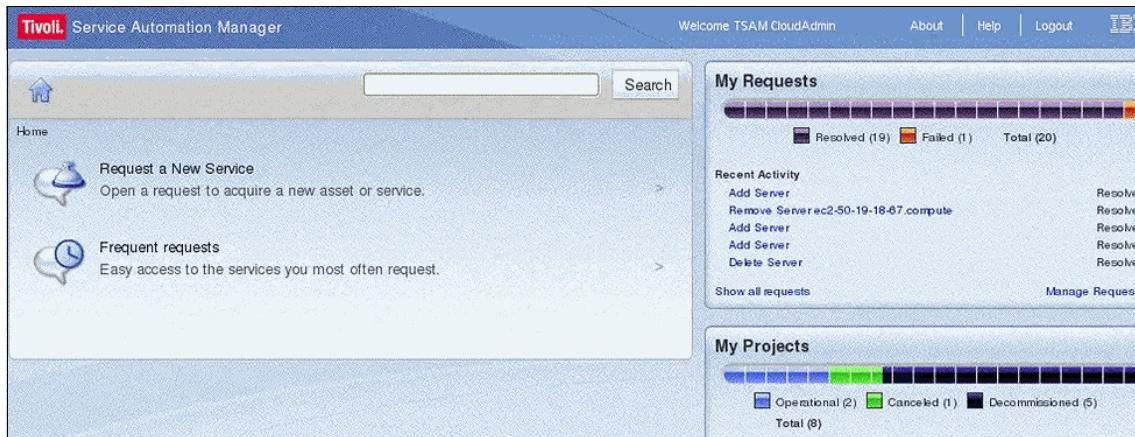


Figure 7-4 Hybrid Cloud Project successfully deployed

Next, we see the resources that are part of the project are being monitored (Figure 7-5). Initially one of the two compute resources that are part of this cloud service is loaded to 100 percent. The other is idle.

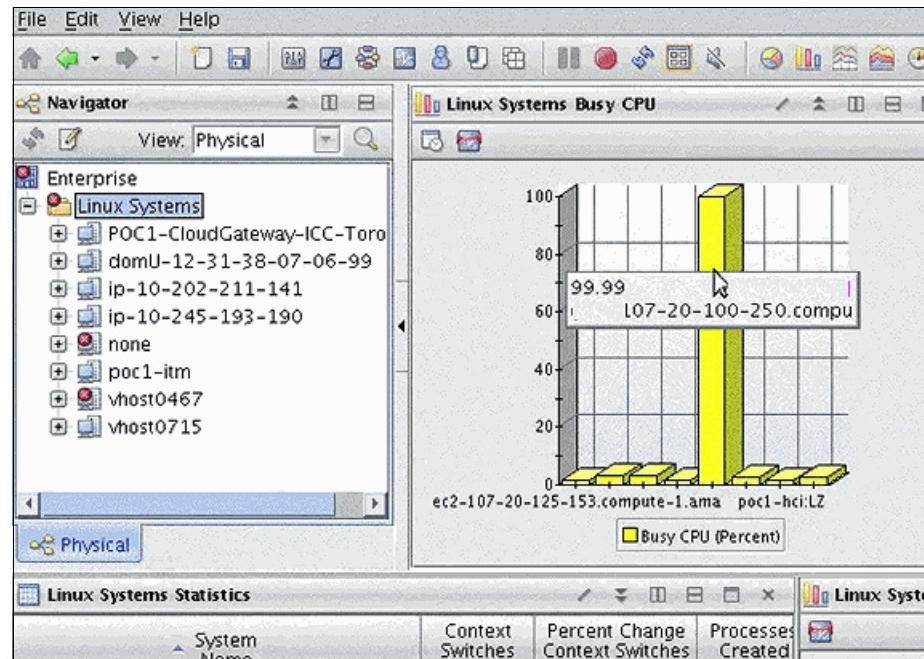


Figure 7-5 CPU utilization is 100 percent on one node

Next we login to the event correlation service in Netcool to see the details of the cloud service performance (Figure 7-6).

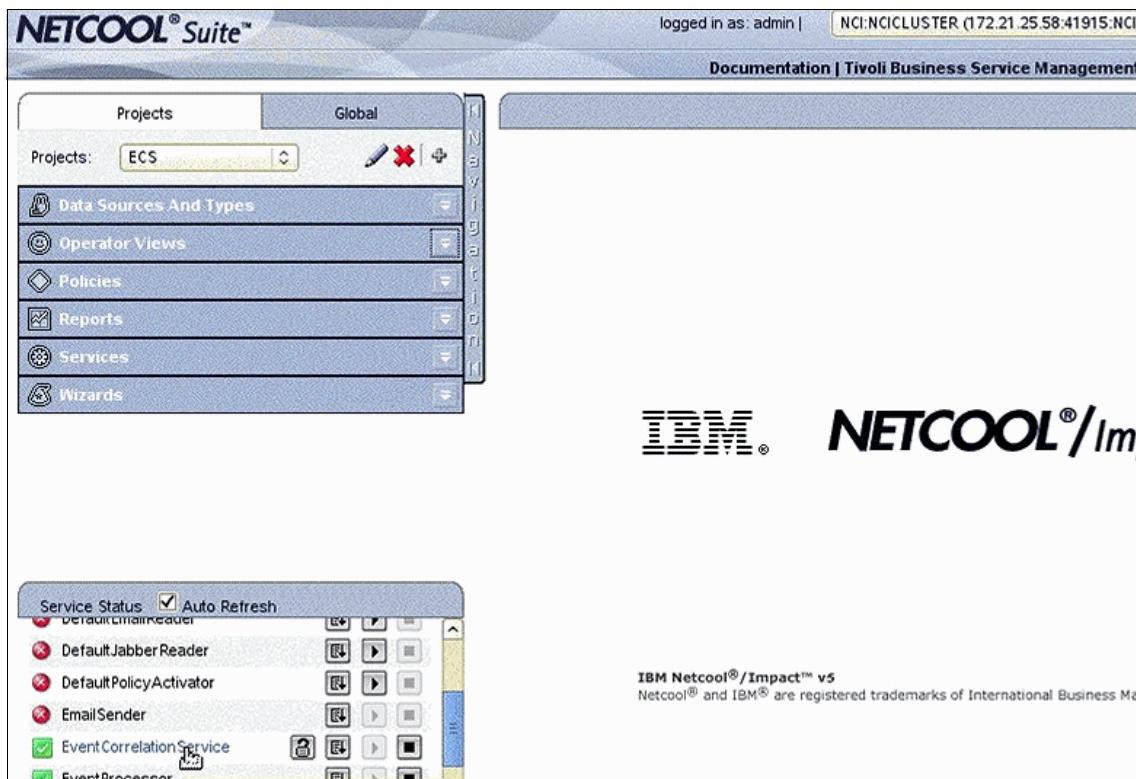


Figure 7-6 Event Correlation Service Console

In the event correlation service log (Figure 7-7) we see that the cloud service is only 50 percent utilized. This is true because one node is at 100 percent and the other is idle.

```
INFO Status: 2 of 2 hosts have repli
.com). 12 samples, 50.88% CPU. Threshold
INFO Quorum (>= 80%) is met, checkin
INFO Synchronizing project list from
INFO Probing data for hosts -107

INFO Status: 2 of 2 hosts have repli
.com). 12 samples, 50.89% CPU. Threshold
INFO Quorum (>= 80%) is met, checkin
INFO Synchronizing project list from
INFO Probing data for hosts -107

INFO Status: 2 of 2 hosts have repli
.com). 12 samples, 50.88% CPU. Threshold
INFO Quorum (>= 80%) is met, checkin
```

Figure 7-7 Event Correlation Service log

Next we add load to the second compute resource. This simulates what would happen during peak retail season. As you can see in Tivoli Monitoring (Figure 7-8) now both nodes are at 100 percent CPU utilization.

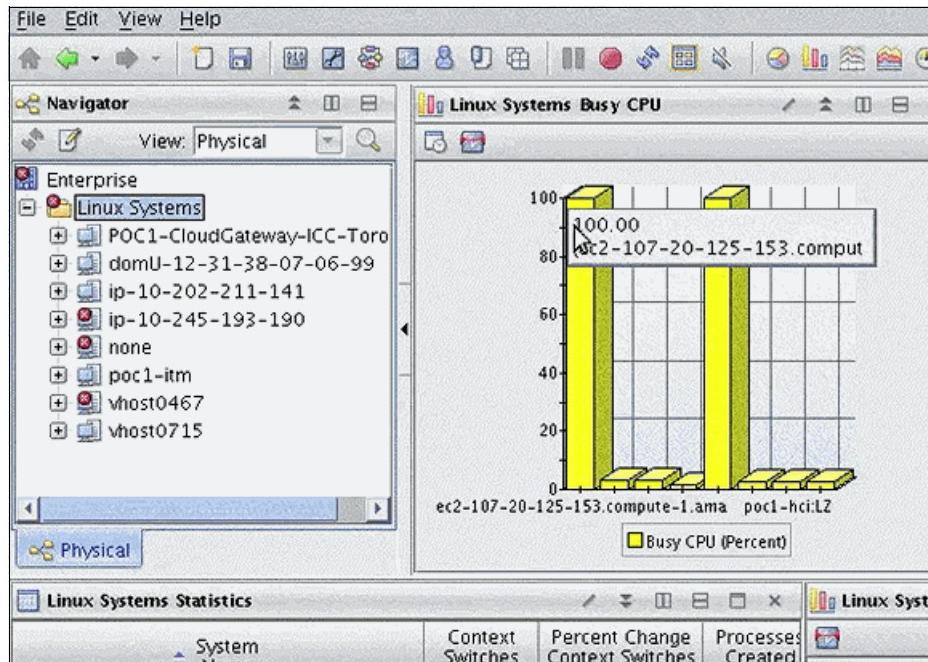


Figure 7-8 CPU Utilization on second node is increased to 100 percent

Next we login to the Service Request Manager to see if the alert has triggered a service request (Figure 7-9).

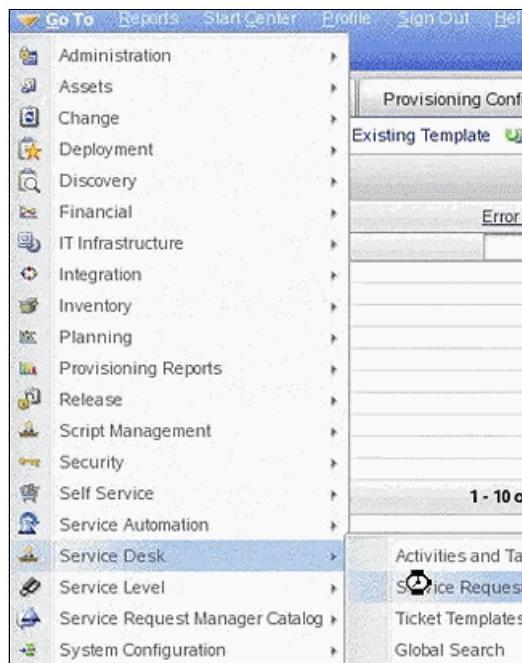


Figure 7-9 Service Request Manager Console

We see that a Workload Overload service request has been generated (Figure 7-10).

| Service Request | Summary              | Reported By    | Internal Priority | Status |
|-----------------|----------------------|----------------|-------------------|--------|
| 1045            | Add Server           | maxadmin       |                   | INPRO  |
| 1044            | Workload Overflow: 7 | ECS Dr. Watson | 1                 | RESOL  |
| 1043            | Add S                | maxadmin       |                   | RESOL  |

Figure 7-10 Workload Service Request created in Service Request System

Next, we look at the details of the service request and see that an automated Workload Overload ticket has been generated and the details of the situation are automatically populated (Figure 7-11).

The screenshot shows a web-based application titled "Service Requests". The main title bar says "Service Requests" and there is a "+" button in the top right corner. Below the title bar, there is a toolbar with icons for "Find" and "Save".

The main content area is divided into sections:

- User Information**:
  - Reported By: [redacted] >>
  - Name: ECS Dr. Watson
  - Phone: [redacted]
  - E-mail: [redacted]
- Service Request Details**:
  - Summary: Workload Overflow: 7
  - Details: Workload has increased significantly for the workload will remain very high for Compute and Storage resources should ECS1316194466382

Figure 7-11 Service Request details

We see that based on our business management policy an automatic ticket has been created to add another compute resource to the cloud service to enforce our policy (Figure 7-12).

The screenshot shows a service request creation form. At the top, there are fields for 'Owner Group' (with a dropdown menu), 'Status' (set to 'INPROG'), and 'Attachments'. Below that is a 'Created By' field containing 'MAXADMIN'. The next section is for 'Contact Person' information, including fields for 'Name' (with a dropdown menu), 'Phone' (with a search icon), and 'E-mail' (with a search icon). At the bottom, there are fields for 'Classification' (set to 'PMRDP\_SR\_ADD\_PUBLIC\_CLOUD\_SERVER'), 'Classification Path' (set to 'PMRDP\_SR\_VS\_MTG \ PMRDP\_SR\_MODIFY\_F'), and 'Class Description' (set to 'Add Servers in Public Cloud').

Figure 7-12 Service Request created for cloud resources to be allocated

In the TSAM interface, we now see the additional resource is in the process of being created (Figure 7-13).

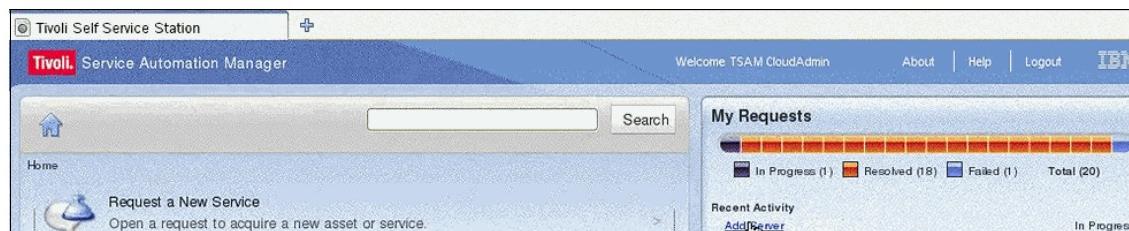
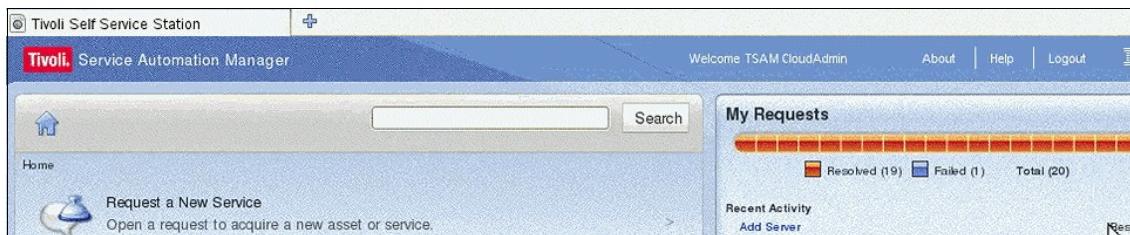


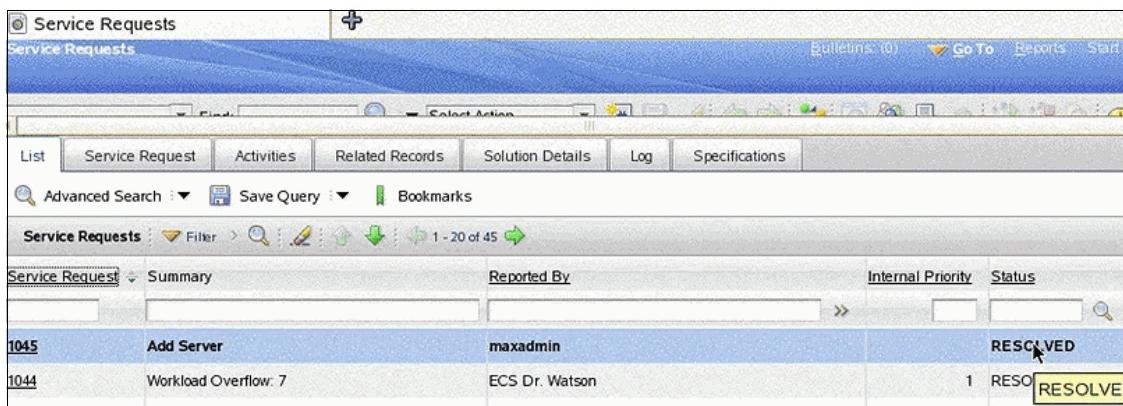
Figure 7-13 Additional compute resource in the process of being allocated

Then, we see that the additional compute resource has been successfully allocated through the TSAM interface (Figure 7-14).



*Figure 7-14 Add New Server has been resolved*

Next, we look in the Service Request system and see that the same status that the additional compute resource has been successfully added (Figure 7-15).



*Figure 7-15 Service Request shows as resolved*

Next, we look in Tivoli Monitoring and see the third node as part of the cloud service and the other two compute nodes still running at 100 percent (Figure 7-16).

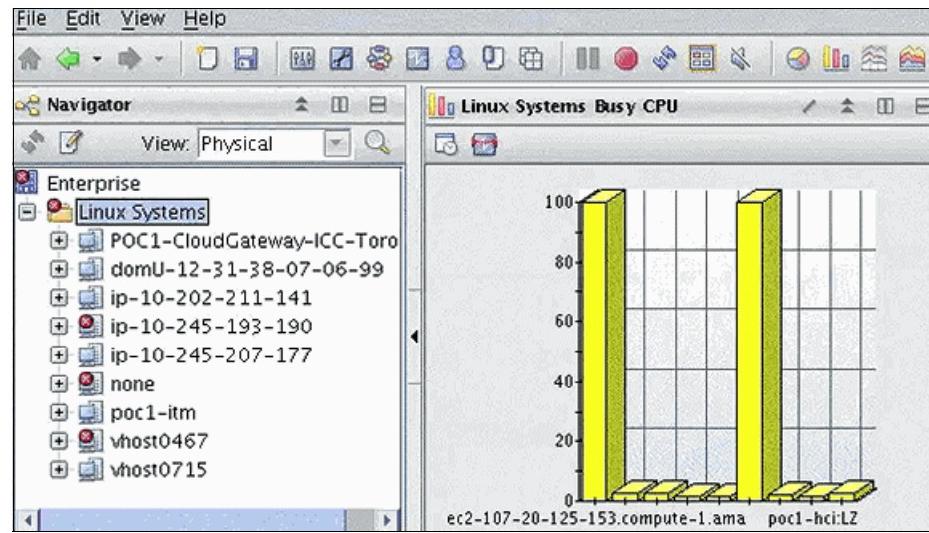


Figure 7-16 Tivoli Monitoring Console showing all three compute resources

Finally we look at the project details in TSAM and see that there are three compute resources are part of the project (Figure 7-17).



Figure 7-17 TSAM project details

### 7.3.2 Resources decommissioned from the cloud project to support business policy

In this section we walk through the Workload Governance and Management scenario after peak retail season is over. We see that CPU utilization on the second compute resource returns to idle and the business policy of keeping CPU utilization between 42 and 75 percent is violated. We then demonstrate the process as the system enforces the policy by decommissioning the unneeded compute resource.

First we see in Tivoli Monitoring that overall utilization for the cloud service has dropped below 42 percent because two of the nodes are not being utilized (Figure 7-18).

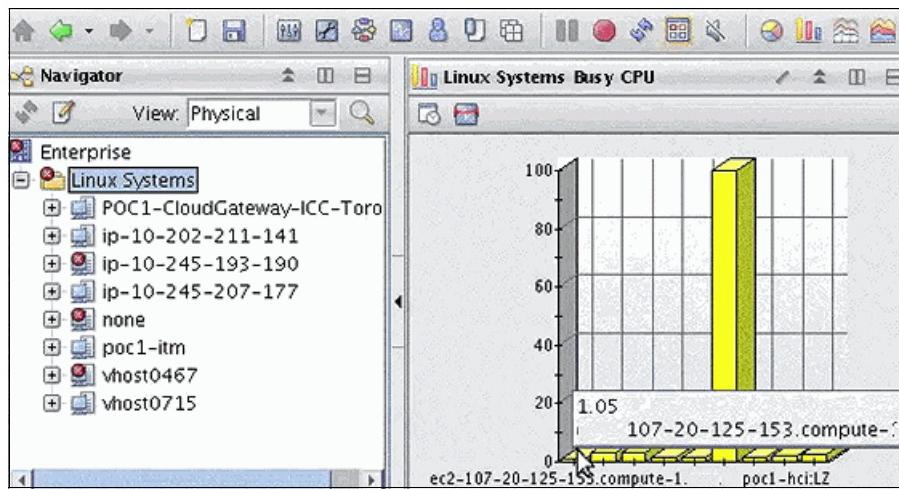


Figure 7-18 Tivoli Monitoring CPU utilization of the nodes that make up the cloud service

Next, we look in the correlation service log and see that the underutilization scenario has been triggered (Figure 7-19).

```
INFO Status: 3 of 3 hosts have replicated data.
INFO Quorum (>= 80%) is met, checking for underutilization.
INFO Synchronizing primary/secondary list ...
INFO Probing data for hosts {ec2-107-20-125-153.compute-1.ip-10-245-193-190.ip-10-245-207-177}
INFO Status: 3 of 3 hosts have replicated data.
INFO Quorum (>= 80%) is met, checking for underutilization.
INFO CPU is 34.46%, and is below threshold.
INFO Deallocating some capacity ...
```

Figure 7-19 Event correlation service Log

Next, we look at the service request to see that a service request has been generated for the underutilized resources (Figure 7-20).

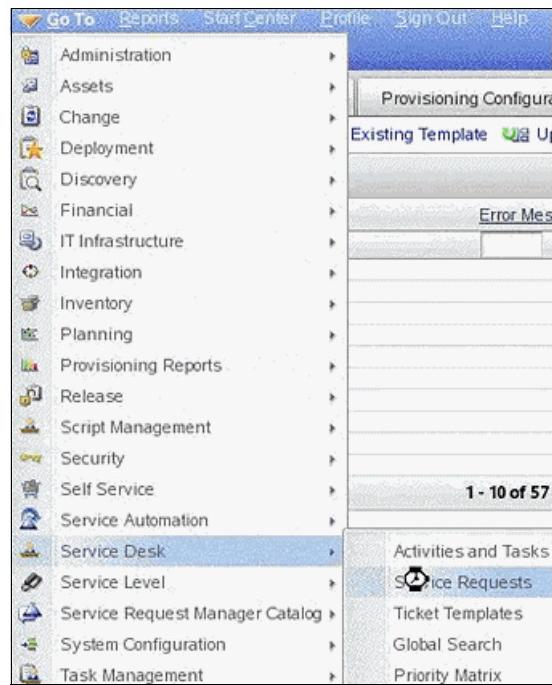


Figure 7-20 Going to Service Details from the Go-To Menu

Next, we see the service request for the underused resource is resolved. and look at the details (Figure 7-21).

The screenshot shows the 'Service Requests' list view. The top navigation bar includes 'Service Requests', 'Go To', and 'Reports'. The main interface has tabs for 'List', 'Service Request', 'Activities', 'Related Records', 'Solution Details', 'Log', and 'Specifications'. Below these are search and filter tools. The main area displays a table of service requests:

| Service Request | Summary              | Reported By    | Internal Priority | Status   |
|-----------------|----------------------|----------------|-------------------|----------|
| 1051            | Delete Server        | maxadmin       |                   | INPROG   |
| 1050            | Workload Underuse: 7 | ECS Dr. Watson | 1                 | RESOLVED |
| 1040            | Delete               | maxadmin       |                   | RESOLVED |

Figure 7-21 Underused resources service request

Figure 7-22 shows the Service Request ticket details.

The screenshot shows a software interface for managing service requests. At the top, there's a header bar with the title 'Service Requests' and a '+' button. Below the header, there are tabs: 'List', 'Service Request' (which is selected), 'Related Records', and 'Solution'. The main content area is divided into sections:

- Service Request:** Number 1050, External ID ECS1316200765749
- User Information:** Reported By (dropdown menu), Name: ECS Dr Watson, Phone (dropdown menu), E-mail (dropdown menu).
- Service Request Details:** Summary: Workload Underuse. 7, Details: Workload has dropped significantly. Workload will remain low for the next.

Figure 7-22 Service request ticket details

Next we see a Service Request to the provisioning service is generated to de-provision the unneeded resources (Figure 7-23).

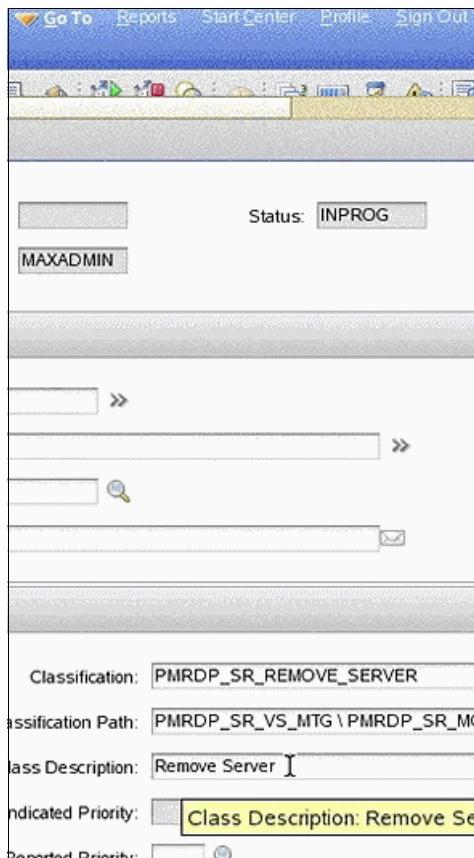


Figure 7-23 Ticket created for provisioning service to decommission the unneeded resources

The Self Service Interface in TSAM shows that the decommissioning of the resource is in progress (Figure 7-24).

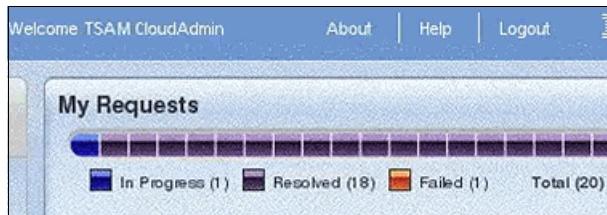


Figure 7-24 Resource decommissioning

We see that the resource has been successfully decommissioned and the utilization of the cloud service is back to within the CPU utilization policy (Figure 7-25).

```
INFO Quorum (>= 80%) is
INFO CPU is 34.46%, and
INFO Deallocating some c
INFO Request sent and wi
INFO Synchronizing projec
INFO Probing data for ho
INFO A host has been remo
INFO Status: 2 of 2 host
.com). 12 samples, 50.9% CP
INFO Quorum (>= 80%) is
```

Figure 7-25 Event Correlation Service log

Now, the ticket in the Self Service Interface has been changed to Resolved status (Figure 7-26).



Figure 7-26 Self Service Interface shows status as resolved

Looking at the details of the project, we see that there are only two nodes in the project and that the third node has been decommissioned (Figure 7-27).

The screenshot shows a web-based interface for Tivoli Self Service Station. The title bar says "Tivoli Self Service Station". A Firefox warning message is displayed: "Do you want Firefox to remember the password for \"PMRDPCAUSR\" on https://poc1-tsam:9443?". The main content area is titled "Project: HYB". It displays "Details for the project and its servers". Under "Project Details", the following information is shown:

|                     |                                     |
|---------------------|-------------------------------------|
| Project Name        | HYB                                 |
| Project Description | (HYB) Governance of Amazon machines |
| Project Type        | RDP                                 |
| Start Date          | 9/15/11 4:53 PM                     |
| End Date            | Indefinite                          |
| Team Access         | ISDMTM1                             |
| Total Server(s)     | 2                                   |
| Active Server(s)    | 2                                   |

Below this, under "Servers", there is a table with the following data:

| Select                           | Server Name               | Management Name | Hypervisor | Memory (%) | CPU (%) | Disk (%) |
|----------------------------------|---------------------------|-----------------|------------|------------|---------|----------|
| <input checked="" type="radio"/> | t-107-20-100-250.compute- |                 |            |            |         |          |
| <input type="radio"/>            | t-107-20-125-153.compute- |                 | No Monitor |            |         |          |

Figure 7-27 Project Details

Finally, we look at the Tivoli Monitoring portal and see that the third node is offline because it has been decommissioned (Figure 7-28).

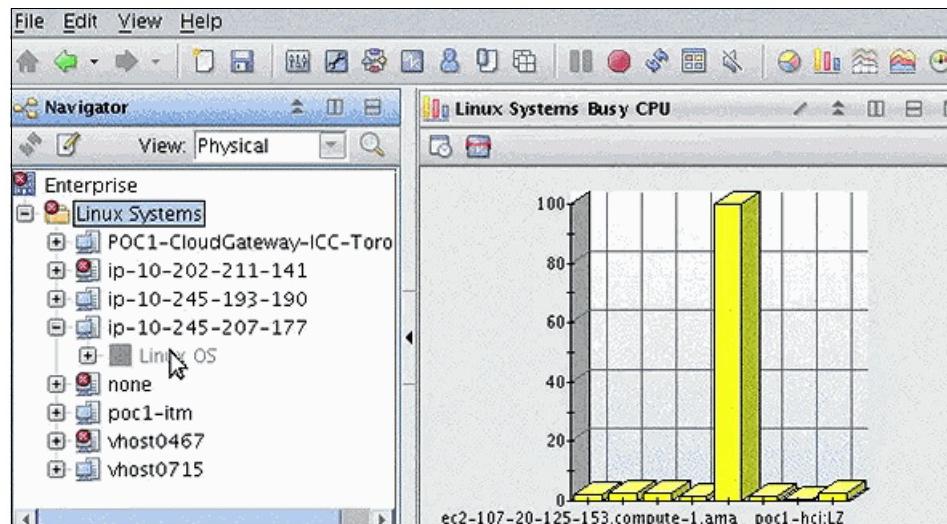


Figure 7-28 Tivoli Monitoring showing the third node is offline

## 7.4 Lessons learned

When working through the Workload Governance and Management scenario, we learned several lessons that make the scenario easier to implement and use.

First, be sure to understand the network connectivity and firewall requirements between the enterprise and the public cloud environment. We ran into this issue in the lab where we had used two layers of network translation and did not forward all the proper ports. This caused the environment to not work until we got the issue resolved.

We also found that all workload governance and management scenario projects must start with "HYB" in the project name. This is because the Event Correlation Service uses this prefix for all its actions. This was an important lesson learned during the usage of this scenario because for the Provisioning and Monitoring scenarios, this was not necessary. This became an issue when we started trying to build on top of the projects that we used for those scenarios, only to realize that the Workload Governance scenario required the HYB prefix for the other components of the solution to work.

## 7.5 Summary

In this chapter we covered the ability to add resources and de-provision resources based on business policy. This allows a company to use both their on-premise resource and public cloud resources effectively to maximize their compute resource investment.

In addition, we explored the ability of a hybrid cloud to enforce policy based on business policy. This enforcement was useful in our use case to ensure that customers were only paying for compute resources they were actually using, while handling seasonal workload without impacting the business.

# Related publications

The publications listed in this section are considered particularly suitable for a more detailed discussion of the topics covered in this book.

## IBM Redbooks

The following IBM Redbooks publications provide additional information about the topic in this document. Note that some publications referenced in this list might be available in softcopy only.

- ▶ *Connect Cloud and On-premise Applications Using IBM WebSphere Cast Iron Integration*, REDP-4674
- ▶ *Strategic Overview of WebSphere Appliances*, REDP-4790
- ▶ *Getting Started with IBM WebSphere Cast Iron Cloud Integration*, SG24-8004
- ▶ *Integration Guide for IBM Tivoli Netcool/OMNibus, IBM Tivoli Network Manager, and IBM Tivoli Netcool Configuration Manager*, SG24-7893
- ▶ *IBM Tivoli Provisioning Manager V7.1.1: Deployment and IBM Service Management Integration Guide*, SG24-7773

You can search for, view, download or order these documents and other Redbooks, Redpapers, Web Docs, draft and additional materials, at the following website:

[ibm.com/redbooks](http://ibm.com/redbooks)

## Online resources

These websites are also relevant as further information sources:

- ▶ WebSphere Cast Iron Cloud Integration home page  
<http://www-01.ibm.com/software/integration/cast-iron-cloud-integration/>
- ▶ WebSphere Cast Iron Cloud Information Center  
<http://publib.boulder.ibm.com/infocenter/wci/v6r1m0/index.jsp>

- ▶ Service Management Extensions for Hybrid Cloud  
<https://www-304.ibm.com/software/brandcatalog/ismlibrary/details?catalog.label=1TW10TS0D>
- ▶ Integrated Service Management Library  
<https://www-304.ibm.com/software/brandcatalog/ismlibrary/>
- ▶ Tivoli Service Request Manager for Service Providers  
<http://www-01.ibm.com/software/tivoli/products/service-request-manager-sp/index.html>
- ▶ Tivoli Service Request Manager for Service Providers Information Center  
[http://pic.dhe.ibm.com/infocenter/tivihelp/v34r1/index.jsp?topic=/com.ibm.srm.doc\\_721/srm\\_welcome\\_sp.html](http://pic.dhe.ibm.com/infocenter/tivihelp/v34r1/index.jsp?topic=/com.ibm.srm.doc_721/srm_welcome_sp.html)
- ▶ Tivoli Service Automation Manager home page  
<http://www-01.ibm.com/software/tivoli/products/service-auto-mgr/>
- ▶ Tivoli Service Automation Manager Information Center  
[http://publib.boulder.ibm.com/infocenter/tivihelp/v10r1/topic/com.ibm.tsam\\_7.2.3.doc/ic-homepage.html](http://publib.boulder.ibm.com/infocenter/tivihelp/v10r1/topic/com.ibm.tsam_7.2.3.doc/ic-homepage.html)
- ▶ IBM SmartCloud Cost Management (Tivoli Usage and Accounting Manager)  
<http://www-01.ibm.com/software/tivoli/products/usage-accounting/>
- ▶ IBM SmartCloud Cost Management (Tivoli Usage and Accounting Manager) Information Center  
[http://pic.dhe.ibm.com/infocenter/tivihelp/v3r1/topic/com.ibm.sccm.doc\\_2.1/welcome.htm](http://pic.dhe.ibm.com/infocenter/tivihelp/v3r1/topic/com.ibm.sccm.doc_2.1/welcome.htm)
- ▶ IBM SmartCloud Control Desk  
<http://www-01.ibm.com/software/tivoli/products/smartcloud-controldesk/>
- ▶ IBM SmartCloud Control Desk Information Center  
<http://pic.dhe.ibm.com/infocenter/tivihelp/v58r1/index.jsp>
- ▶ Tivoli Application Dependency Discovery Manager  
<http://www-01.ibm.com/software/tivoli/products/taddm/>
- ▶ Tivoli Application Dependency Discovery Manager Information Center  
[http://pic.dhe.ibm.com/infocenter/tivihelp/v46r1/index.jsp?topic=%2Fcom.ibm.taddm.doc\\_721fp3%2Fwelcome\\_page%2Fwelcome.html](http://pic.dhe.ibm.com/infocenter/tivihelp/v46r1/index.jsp?topic=%2Fcom.ibm.taddm.doc_721fp3%2Fwelcome_page%2Fwelcome.html)
- ▶ Tivoli Provisioning Manager  
<http://www-01.ibm.com/software/tivoli/products/prov-mgr/>

- ▶ Tivoli Provisioning Manager Information Center  
<http://pic.dhe.ibm.com/infocenter/tivihelp/v45r1/topic/com.ibm.tivoli.tpm.doc/welcome/ic-homepage.html>
- ▶ Tivoli Service Level Advisor  
<http://www-01.ibm.com/software/tivoli/products/service-level-advisor/>
- ▶ Tivoli Netcool Service Quality Manager  
<http://www-01.ibm.com/software/tivoli/products/netcool-service-quality-mgr/>
- ▶ IBM System Director VMControl  
<http://www-03.ibm.com/systems/software/director/vmcontrol/index.html>
- ▶ IBM Tivoli Storage Productivity Center  
<http://www-03.ibm.com/systems/storage/software/center/>
- ▶ IBM Tivoli Monitoring Information Center  
[http://pic.dhe.ibm.com/infocenter/tivihelp/v15r1/topic/com.ibm.itm.doc\\_6.2.3fp1/welcome.htm](http://pic.dhe.ibm.com/infocenter/tivihelp/v15r1/topic/com.ibm.itm.doc_6.2.3fp1/welcome.htm)
- ▶ Social Business IBM SmartCloud  
<https://www.lotuslive.com/en/>
- ▶ IBM SmartCloud Enterprise  
<https://www-147.ibm.com/cloud/enterprise/dashboard>

## Help from IBM

IBM Support and downloads

[ibm.com/support](http://ibm.com/support)

IBM Global Services

[ibm.com/services](http://ibm.com/services)





## Hybrid Cloud Integration and Monitoring with IBM WebSphere Cast Iron

(0.2"spine)  
0.17" <-> 0.473"  
90<->249 pages





# Hybrid Cloud Integration and Monitoring with IBM WebSphere Cast Iron



Redbooks®

## Integrate cloud systems with on-premise systems

## Use a unified interface for on- and off-premise cloud systems

## Allocate cloud resources based on business needs

Cloud computing is fast emerging as a new consumption and delivery model for IT solutions. Two distinct cloud patterns have evolved. Private clouds are custom cloud applications that are on premise to companies. Public clouds are owned and managed by a service provider and are delivered on a pay-per-use basis. Although the debate between the use of private versus public cloud continues, the industry is moving ever faster toward a hybrid solution. Hybrid clouds address the integration of traditional IT environments with one or more clouds, private and public.

IBM WebSphere Cast Iron along with the Tivoli Service Management Extensions for Hybrid Cloud supports the integration of private and public clouds. The Service Management Extensions for Hybrid Cloud supports the following scenarios:

- ▶ Provisioning hybrid cloud resources
- ▶ Monitoring hybrid cloud resources
- ▶ Governing and managing workloads that use hybrid cloud resources
- ▶ Provisioning users from LDAP to LotusLive

This IBM Redbooks publication is intended for application integrators, integration designers, and administrators evaluating or using IBM WebSphere Cast Iron. In addition, executives, business leaders, and architects who are looking for a way to integrate cloud applications with their on-premise applications are shown how WebSphere Cast Iron can help to resolve their integration challenges. The book helps you gain an understanding of hybrid cloud use cases, and explains how to integrate cloud and on-premise applications quickly and simply.

**INTERNATIONAL  
TECHNICAL  
SUPPORT  
ORGANIZATION**

**BUILDING TECHNICAL  
INFORMATION BASED ON  
PRACTICAL EXPERIENCE**

IBM Redbooks are developed by the IBM International Technical Support Organization. Experts from IBM, Customers and Partners from around the world create timely technical information based on realistic scenarios. Specific recommendations are provided to help you implement IT solutions more effectively in your environment.

**For more information:  
[ibm.com/redbooks](http://ibm.com/redbooks)**